arXiv:1703.07710v1 [cs.LG] 22 Mar 2017

UBEV – A More Practical Algorithm for Episodic RL with Near-Optimal PAC and Regret Guarantees

Christoph Dann [1] **Tor Lattimore** [2] **Emma Brunskill** [1 3]

Abstract

We present UBEV, a simple and efficient reinforcement learning algorithm for fixed-horizon episodic Markov decision processes. The main contribution is a proof that UBEV enjoys a sample-complexity bound that holds for all accuracy levels simultaneously with high probability, and matches the lower bound except for logarithmic terms and one factor of the horizon. A consequence of the fact that our sample-complexity bound holds for all accuracy levels is that the new algorithm achieves a sub-linear regret of $\tilde{O}(\sqrt{SAT})$, which is the first time the dependence on the size of the state space has provably appeared inside the square root. A brief empirical evaluation shows that UBEV is practically superior to existing algorithms with known sample-complexity guarantees.

1. Introduction

Sample efficient reinforcement learning is an important challenge in sequential decision making. While naive approaches to exploration like ε-greedy are sometimes effective in challenging tasks (Mnih et al., 2015), they usually perform poorly when data is scarce or in domains that require directed exploration. For this reason we approach the problem of exploration with the objective of designing robust algorithms that explore more carefully than heuristics.

We introduce a simple new algorithm for provably sample efficient reinforcement learning in tabular episodic Markov decision processes (MDP). While there are clear limitations to tabular representations, it is perhaps the simplest reinforcement learning setting where actions influence states (not bandits) and we anticipate advances in understanding efficient learning in this setting may also bring later insights for larger domains, as has been observed recently, for example, by Bellemare et al. (2016).

[1]Carnegie Mellon University, Pittsburgh, USA [2]Independent Researcher [3]Stanford University, Palo Alto, USA. Correspondence to: Christoph Dann <cdann@cdann.net>.

We measure the performance of our algorithm in the PAC-RL framework (Kearns & Singh, 2002; Kakade, 2003). A reinforcement learning algorithm is PAC-RL if with high probability it behaves near-optimally for all but at most a polynomial number of episodes, where the polynomial (sometimes called the sample-complexity) depends on the failure probability, the definition of 'near-optimal' and on the size/structure of the MDP. Of the existing algorithms with PAC guarantees, UCFH and MoRMax have the smallest sample-complexity (see Table 1). We prove a lower bound that is linear in S and cubic in H, which we conjecture is optimal and suggests that UCFH is suboptimal in terms of the state-space S and MoRMax is suboptimal in terms of the horizon H. Equally important, existing analyses generally require that the relevant algorithm eventually stop updating its policy (see, for example, (Strehl & Littman, 2004; Strehl et al., 2009; Dann & Brunskill, 2015; Szita & Szepesvári, 2010)) and often discard large amounts of data. Both are counter to the intuition that more experience should yield a better (or at least not worse) policy.

MBIE†	Del. Q-Lear.†	MoRMax†	UCFH*	**UBEV**
$\frac{H^6S^2A}{\varepsilon^2}$	$\frac{H^8SA}{\varepsilon^3}$	$\frac{H^6SA}{\varepsilon^2}$	$\frac{H^3S^2A}{\varepsilon^2}$	$\frac{H^4SA}{\varepsilon^2}$

S = size of state space, A = size of action space, H = episode length, ε = desired accuracy. All bounds hold up to poly-logarithmic terms.

† The sample complexity of algorithms were originally analysed as time-step bound in the discounted setting with discount factor $\gamma \in (0,1)$. However, the results and analysis can be translated to bounds on the number of ε-suboptimal episodes in the finite-horizon episodic case by substituting the effective horizon $1/(1-\gamma)$ with the episode length H.

$*$ There is an additional H in the sample complexity for our setting due to the potentially time-dependent dynamics.

Table 1: Summary of PAC bounds

Contributions To address these issues and some of the open gaps in the theory, we introduce a simple and efficient algorithm that uses a new approach to Upper Bound the Expected Next Value (UBEV) when computing an optimistic policy to drive exploration. UBEV is a PAC-RL algorithm

with a sample complexity that for sufficiently small accuracy ε is optimal except for one factor of the horizon. Importantly, its sample complexity scales linearly with the size of the state space, which is a significant improvement over the quadratic dependence of horizon-optimal algorithms (Lattimore & Hutter, 2014; Dann & Brunskill, 2015). We also introduce a variant of the PAC optimality criteria called δ-*uniform PAC*. Recall an algorithm is (ε, δ)-PAC if

$$\mathbb{P}\left(N_\varepsilon > F\left(S, A, H, \frac{1}{\varepsilon}, \log\frac{1}{\delta}\right)\right) \le \delta\,,$$

where N_ε is the (random) number of episodes where the algorithm is following an ε-suboptimal policy and F a is polynomial (preferably with small degree). We call an algorithm δ-*uniform-PAC* if

$$\mathbb{P}\left(\text{exists } \varepsilon > 0 : N_\varepsilon > F\left(S, A, H, \frac{1}{\varepsilon}, \log\frac{1}{\delta}\right)\right) \le \delta\,.$$

The difference is that a δ-uniform-PAC algorithm does not depend on a desired accuracy and should, with high probability, make only a polynomial number of ε-errors for all $\varepsilon > 0$ simultaneously. An implication of the new definition is that with probability at least $1-\delta$, a δ-uniform-PAC algorithm will converge to the optimal policy. This is not generally true for (ε, δ)-PAC algorithms, which usually cease exploration once the dynamics are learned to a sufficient accuracy. Another advantage is that δ-uniform-PAC algorithms can be shown to enjoy sub-linear regret, which for newly proposed algorithm scales as $\tilde{O}(H^2\sqrt{SAT})$. This is the first time the dependence on S has been moved inside the square root, which matches the known lower bound regret dependence on S (Auer et al., 2010). Finally, UBEV never stops learning and updates its policy after every episode, which in our simulations provides a benefit over previous PAC-RL algorithms.

2. Background and Setting

We consider episodic fixed-horizon MDPs with time-dependent dynamics, which can be formalized as a tuple $M = (\mathcal{S}, \mathcal{A}, p_R, P, p_0, H)$. The statespace $\mathcal{S}$ and the actionspace $\mathcal{A}$ are finite sets with cardinality S and A. The agent interacts with the MDP in episodes of H time steps. At the beginning of each time-step $t \in [H]$ the agent observes a state s_t and chooses an action a_t based on a policy π that may depend on the within-episode time step ($a_t = \pi(s_t, t)$). The next state is sampled from the tth transition kernel $s_{t+1} \sim P(\cdot|s_t, a_t, t)$ and the initial state from $s_1 \sim p_0$. The agent then receives a reward drawn from a distribution $p_R(s_t, a_t, t)$ which can depend on s_t, a_t and t with mean $r(s_t, a_t, t)$ determined by the reward function. The reward distribution p_R is supported on $[0, 1]$.[1] The value function from time step t for policy π is defined as

$$\begin{aligned} V_t^\pi(s) &:= \mathbb{E}\left[\sum_{i=t}^{H} r(s_i, a_i, i) | s_t = s\right] \\ &= \sum_{s' \in \mathcal{S}} P(s'|s, \pi(s,i), i) V_{i+1}^\pi(s') + r(s, \pi(s,i), i)\,. \end{aligned}$$

and the optimal value function is denoted by V_t^*. In any fixed episode, the quality of a policy π is evaluated by the *total expected reward* or *return*

$$\mathbb{E}\left[\sum_{i=t}^{H} r(s_i, a_i, i)\middle|\pi\right] = p_0^\top V_1^\pi$$

which is compared to the *optimal return* $p_0^\top V_1^*$. As previously mentioned, the objective of this article is to design a (random) sequence of policies $\pi_1, \pi_2, \ldots$ (one for each episode) so with probability at least $1-\delta$, the number of episodes k when $p_0^\top V_1^* - \varepsilon > p_0^\top V_1^{\pi_k}$ is as small as possible for all $\varepsilon > 0$ simultaneously.[2]

Notation Given an integer x, we use the short hand notation $[x] = \{1, 2, \ldots, x\}$. The function $\mathrm{llnp}(x) = \ln\ln\max\{e, x\}$ is a truncated version of the nested logarithm $\ln(\ln(x))$ and the range of a vector v is $\mathrm{rng}\, v = \max v - \min v$. With slight abuse of notation we denote by $P(s, a, t) \in [0, 1]^S$ the vector with probabilities $P(\cdot|s, a, t)$. We will reason about the true MDP M, an empirical MDP $\hat{M}_k$ and an optimistic MDP $\tilde{M}_k$ which are identical except for their transition probabilities P, $\hat{P}_k$ and $\tilde{P}_k$ and reward functions r, $\hat{r}_k$ and $\tilde{r}_k$. We will provide more details later. We introduce the notation explicitly only for M but the quantities carry over with tildes or hats.

3. The UBEV Algorithm

The pseudo-code for UBEV is given in Algorithm 1. In each episode it follows an optimistic policy π_k that is computed by backwards induction using a carefully chosen confidence interval on the transition probabilities in each state. We show in Lemma B.1 in the appendix that the policy update in Lines 4–10 finds an optimal solution to the following optimization problem which is the problem of finding the optimal policy but where the MDP parameters P', r' are also optimization variables:

$$\max_{P', r', V', \pi'} \ \mathbb{E}_{s\sim p_0}[V_1'(s)] \tag{1}$$

[1] The reward may be allowed to depend on the next-state with no further effort in the proofs. The boundedness assumption could be replaced by the assumption of subgaussian noise with known subgaussian parameter.

[2] The number of ε-errors is here defined as $N_\varepsilon = \sum_{k=1}^{\infty} \mathbb{I}\{p_0^\top (V_1^* - V_1^{\pi_k}) > \varepsilon\}$ where $\mathbb{I}\{\cdot\}$ is the indicator function.

subject to the constraints that for all $s \in \mathcal{S}, a \in \mathcal{A}, t \in [H]$,

(a) $V'_{H+1} = 0, \quad P'(s,a,t) \in \Delta_S, \quad r'(s,a,t) \in [0,1]$.

(b) $V'_t(s) = r(s,\pi'(s,t),t) + P'(s,\pi'(s,t),t)^\top V'_{t+1}$.

(c) $|[(P' - \hat{P}_k)(s,a,t)]^\top V'_{t+1}| \le \phi(s,a,t)\,\mathrm{rng}\, V'_{t+1}$.

(d) $|r'(s,a,t) - \hat{r}_k(s,a,t)| \le \phi(s,a,t)$,

where $(P' - \hat{P}_k)(s,a,t)$ is short for $P'(s,a,t) - \hat{P}_k(s,a,t)$ and

$$\phi(s,a,t) = \sqrt{\frac{2\,\mathrm{llnp}(n(s,a,t)) + \ln(18SAH/\delta)}{n(s,a,t)}}$$

is the width of a confidence bound, $\hat{P}_k(s'|s,a,t) = \frac{m(s',s,a,t)}{n(s,a,t)}$ are the empirical transition probabilities and $\hat{r}_k(s,a,t) = l(s,a,t)/n(s,a,t)$ the empirical immediate rewards (both at the beginning of the kth episode). Our algorithm is conceptually similar to other algorithms based on the optimism principle such as MBIE, UCFH, UCRL2 (Jaksch et al., 2010) or UCRL-γ (Lattimore & Hutter, 2014) but there are several key differences:

- Instead of using confidence intervals over the transition kernel by itself, we incorporate the value function directly into the concentration analysis as in (c) above. Ultimately this saves a factor of S in the sample complexity, but the price is a more difficult analysis. Previously MoRMax(Szita & Szepesvári, 2010) also used the idea of directly bounding the transition and value function, but in a very different algorithm that required discarding data and had a less tight bound.
- Many algorithms update their policy less and less frequently (usually when the number of samples doubles), and only finitely often in total. Instead, we update the policy after every episode, which means that UBEV immediately leverages new observations.
- The policy update of UBEV that occurs once per each episode is computationally efficient. It has the same runtime as value iteration for a fixed MDP $O(S^2AH)$.

Algorithm 1: UBEV (**U**pper **B**ounding the **E**xpected Next State **V**alue) Algorithm

Input: failure tolerance $\delta \in (0,1]$

for $s, s' \in \mathcal{S}, a \in \mathcal{A}, t \in [H]$ **do**
  $n(s,a,t) = l(s,a,t) = m(s',s,a,t) = 0$
  $\tilde{V}_{H+1}(s') := 0$
for $k = 1, 2, 3, \ldots$ **do**
  /* Optimistic planning */
  for $t = H$ **to** 1 **do**
    for $s \in \mathcal{S}$ **do**
      for $a \in \mathcal{A}$ **do**
        $\phi := \sqrt{\frac{2\,\mathrm{llnp}(n(s,a,t)) + \ln(18SAH/\delta)}{n(s,a,t)}}$
        $EV(a) := \min\left\{1, \frac{l(s,a,t)}{\max\{1,n(s,a,t)\}} + \phi\right\}$
        $+ \min\left\{\max \tilde{V}_{t+1}, \frac{m(\cdot,s,a,t)^\top \tilde{V}_{t+1}}{\max\{1,n(s,a,t)\}} + \mathrm{rng}\,\tilde{V}_{t+1}\phi\right\}$
      $\pi_k(s,t) := \arg\max_a EV(a)$
      $\tilde{V}_t(s) := EV(\pi_k(s,t))$
  /* Execute policy for one episode */
  $s_1 \sim p_0$;
  for $t = 1$ **to** H **do**
    $a_t := \pi_k(s_t,t)$
    $r_t \sim p_R(s_t,a_t,t)$ and $s_{t+1} \sim P(s_t,a_t,t)$
    $n(s_t,a_t,t) := n(s_t,a_t,t) + 1$
    $m(s_{t+1},s_t,a_t,t) := m(s_{t+1},s_t,a_t,t) + 1$
    $l(s_t,a_t,t) := l(s_t,a_t,t) + r_t$

4. PAC Analysis

Theorem 1. *Let π_k be the policy of UBEV in the kth episode. Then with probability at least $1-\delta$ for all $\varepsilon > 0$ jointly the number of episodes k where the expected return from the start state is not ε-optimal ($p_0^\top V_1^* - p_0^\top V_1^{\pi_k} > \varepsilon$) is at most*

$$O\left(\frac{SAH^4}{\varepsilon^2}\min\left\{1+\varepsilon S^2 A, S\right\}\mathrm{polylog}\left(A,S,H,\frac{1}{\varepsilon},\frac{1}{\delta}\right)\right).$$

Here $\mathrm{polylog}(x\ldots)$ is a function that can be bounded by a polynomial of logarithm, that is, $\exists k, C : \mathrm{polylog}(x\ldots) \le \ln(x\ldots)^k + C$. Some comments are in order:

- Our PAC bound for UBEV is never worse than $\tilde{O}(S^2AH^4/\varepsilon^2)$, which improves on the similar MBIE algorithm by a factor of H^2 (after adapting the discounted for which MBIE was analysed to our setting). For $\varepsilon < 1/(S^2A)$ our bound has a linear dependence on the size of the state-space and depends on H^4, which is a tighter dependence on the horizon than MoRMax's $\tilde{O}(SAH^6/\varepsilon^2)$, the best sample-complexity bound with linear dependency S so far. Besides the theoretical improvement, UBEV is a simple and efficient algorithm that makes use of all the data (unlike MoRMax).
- Previous sample-complexity guarantees have all been proved for a pre-specified value of ε on which the algorithm depended. In contrast, the new algorithm does not rely on the knowledge of ε and the order of quantifiers in the statement of the theorem means that with high probability the algorithm makes only a polynomial

number of ε-errors for all $\varepsilon > 0$ simultaneously. This immediately implies that with high probability UBEV eventually converges to the optimal policy, which was not true for existing PAC RL algorithms.

- We leverage the fact that our PAC bound holds for all $\varepsilon > 0$ simultaneously to derive an $\tilde{O}(\sqrt{SAT})$ regret guarantee. This makes UBEV the first strategy for which a square root dependence on the size of the state-space has been established. Previous attempts had a linear dependence on this quantity, though a Bayesian regret with square root dependence was recently established (Osband et al., 2016).
- In Sec. 4.3 we provide a lower bound on the sample complexity that shows that if $\varepsilon < 1/(S^2A)$, the PAC bound is tight up to logarithmic factors and a factor of H.

The full proof of Theorem 1 is in the appendix. In order to show that the policy π_k followed by UBEV is close to optimal, we have to assume that the observed state and reward samples satisfy a regularity condition, that is, empirical averages (like $\hat{P}_k, \hat{r}_k$) are reasonably close to their true expected values, satisfying confidence bounds such as $|\hat{r}_k(s,a,t) - r(s,a,t)| \leq \phi(s,a,t)$. Otherwise the observations are not typical for the true MDP and mislead the algorithm.

The first of two main technical components of the proof for Theorem 1 is to show that with probability at least $1-\delta$, this regularity condition is satisfied for all episodes simultaneously. The confidence bounds in the regularity conditions of previous algorithms explicitly depend on the desired accuracy ε of the (ε, δ)-PAC guarantee but ours are independent of ε which is crucial for proving δ-uniform PAC guarantees. Our analysis uses finite-time versions of the law of the iterated logarithm which is why $\mathrm{llnp}(\cdot) \approx \ln\ln(\cdot)$ appears for example in $\phi(s,a,t)$. For details, see Appendix F.Such uniform concentration guarantees have been known for a long time, but only recently have become popular in the ML community. To our knowledge we are the first to use them for PAC RL.

The second part of the proof is to control the number of ε-suboptimal episodes under the assumption that the observations satisfy the regularity condition. This second component is more interesting and is summarised in some detail in the next two subsections.

4.1. Decomposing the Optimality Gap

In the rest of this section, we will decompose the difference between the expected sum of rewards under the optimal policy $p_0^\top V_1^*$ and the expected sum of rewards $p_0^\top V_1^{\pi_k}$ of UBEV's current policy π_k . We call this the optimality gap. Throughout this sketch of our analysis we assume that the observations satisfy the regularity conditions (see above, details in the appendix) and want to prove that the optimality gap is no more than ε on all but a bounded number of episodes.

The true MDP P, r and optimal policy V^*, π_* are a feasible solution of the planning problem in (1) and due to this optimism of UBEV, we can transform the optimality gap into a value difference for the same policy.

$$p_0^\top V_1^* - p_0^\top V_1^{\pi_k} \leq p_0^\top (\tilde{V}_1^{\pi_k} - V_1^{\pi_k}) \,.$$

The right hand side in the expression above is then decomposed via a standard identity (Lemma C.15 in the appendix) as

$$\begin{aligned} &p_0^\top (\tilde{V}_1^{\pi_k} - V_1^{\pi_k}) \\ =&\sum_{t=1}^{H} \sum_{s,a \in \mathcal{S}\times\mathcal{A}} w_{tk}(s,a)((\tilde{P}_k - P)(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k} \\ &+ \sum_{t=1}^{H} \sum_{s,a \in \mathcal{S}\times\mathcal{A}} w_{tk}(s,a)(\tilde{r}_k(s,a,t) - r(s,a,t)), \end{aligned}$$

where $w_{tk}(s,a)$ is the probability that when following policy π_k in the true MDP we encounter $s_t = s$ and $a_t = a$. The quantities $\tilde{P}_k$, $\tilde{r}_k$ are the model parameters of the optimistic MDP before episode k, that is, the optimal values of P' and r' in (1). For the sake of conciseness, we ignore the second term above in the following which can be bounded in the same way as the first. We further decompose the first term as

$$\sum_{\substack{t\in[H] \\ (s,a)\in L_{tk}^c}} w_{tk}(s,a)((\tilde{P}_k - P)(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k} \tag{2}$$

$$+ \sum_{\substack{t\in[H] \\ (s,a)\in L_{tk}}} w_{tk}(s,a)((\tilde{P}_k - \hat{P}_k)(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k} \tag{3}$$

$$+ \sum_{\substack{t\in[H] \\ (s,a)\in L_{tk}}} w_{tk}(s,a)((\hat{P}_k - P)(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k} \tag{4}$$

where

$$L_{tk} = \left\{ (s,a) \in \mathcal{S}\times\mathcal{A} \,:\, w_{tk}(s,a) \geq w_{\min} = \frac{\varepsilon}{3HS^2} \right\}$$

is the set of state-action pairs with non-negligible visitation probability. The value of $w_{\min}$ is chosen so that (2) is bounded by

$$\sum_{\substack{t\in[H] \\ (s,a)\in L_{tk}^c}} w_{\min}((\tilde{P}_k - P)(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k} \leq \frac{\varepsilon}{3} \,.$$

Since $\tilde{V}^{\pi_k}$ is the optimal solution of the planning problem

in Eq. (1), we can bound

$$|((\tilde{P}_k - \hat{P}_k)(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k}| \le \phi_k(s,a,t) \operatorname{rng} \tilde{V}_{t+1} = O\left(\sqrt{\frac{H^2 \ln \frac{\ln n_{tk}(s,a)}{\delta}}{n_{tk}(s,a)}}\right),$$

where $\phi_k(s,a,t)$ is the value of $\phi(s,a,t)$ and $n_{tk}(s,a)$ the value of $n(s,a,t)$ right before episode k. Substituting this into (3) shows that

$$(3) = O\left(\sum_{\substack{\sum_{t\in[H]} \\ s,a\in L_{tk}}} w_{tk}(s,a)\sqrt{\frac{H^2 \ln \frac{\ln n_{tk}(s,a)}{\delta}}{n_{tk}(s,a)}}\right). \quad (5)$$

Recall that $w_{tk}(s,a)$ the probability that state-action-time triple (s,a,t) is encountered in the kth episode, so we should expect that with reasonably high probability,

$$n_{tk}(s,a) \approx \sum_{u<k} w_{tu}(s,a). \quad (6)$$

The right-hand-side of (5) can only be big if $w_{tk}(s,a)$ is reasonably large for (s,a,t) triples where $n_{tk}(s,a)$ is small, which by Equation (6) cannot happen too often. To precisely bound how often this occurs we need a careful argument (Appendix C.3), which is laid out below concisely. Let $\mathcal{I}$ and $\mathcal{K}$ be index sets and for $i,k \in \mathcal{I}\times\mathcal{K}$ let w_{ik} be positive so that $\sum_{i\in\mathcal{I}} w_{ik} \le H$ and $W_{ik} = 1 + \sum_{u<k} w_{ik}$. Then

$$\sum_{k\in\mathcal{K}}\left(\sum_{i\in\mathcal{I}} w_{ik}\sqrt{\frac{1}{W_{ik}}}\right)^2 \le \sum_{k\in\mathcal{K}}\left(\sum_{i\in\mathcal{I}} w_{ik}\right)\sum_{i\in\mathcal{I}}\frac{w_{ik}}{W_{ik}} \lesssim H\sum_{i\in\mathcal{I}} \log W_{i|\mathcal{K}|} \lesssim H|\mathcal{I}|\operatorname{polylog}(\cdot),$$

where the first inequality follows from Cauchy-Schwarz inequality, the second since $\int f'/f d\lambda = \log f$, the third by simply being vague about some logarithmic terms (of course this is dealt with carefully in the appendix). This result can be applied to (5) by letting $\mathcal{I} = \mathcal{S}\times\mathcal{A}\times[H]$ and $w_{ik} = w_{tk}(s,a)$ where $i = (s,a,t)$. Note that by definition of $w_{tk}(s,a)$ it does indeed hold that $\sum_{i\in\mathcal{I}} w_{ik} = H$. We also let $W_{ik} + 1 = n_{tk}(s,a)$. The $+1$ is reasonable on the basis that $n_{tk}(s,a) \ge 1$ except for a very few episodes. Then let $\mathcal{K}$ be the set of episodes where Eq. (5) $\ge \varepsilon/3$. So ignoring $\ln(\ln(\cdot)/\delta)$ in (5) gives

$$\frac{|\mathcal{K}|\varepsilon^2}{9} = \sum_{k\in\mathcal{K}}\frac{\varepsilon^2}{9} \le \sum_{k\in\mathcal{K}} (\text{Eq.5})^2 \lesssim H^3|\mathcal{I}|\operatorname{polylog}(\cdot) = H^4 SA \operatorname{polylog}(A,S,H,1/\varepsilon,1/\delta).$$

Therefore by re-arranging, the number of episodes where (5) is larger than $\varepsilon/3$ is at most

$$O\left(\frac{SAH^4}{\varepsilon^2}\operatorname{polylog}(A,S,H,1/\varepsilon,1/\delta)\right). \quad (7)$$

Note that here we have a linear dependence on S. To bound (4) we argue in a similar way, but now we cannot rely on the definition of the algorithm to control the inner product, as we did for (3). The normal method is to use a concentration analysis, but both $\hat{P}_k$ and $\tilde{V}^{\pi_k}$ are random and dependent, which makes standard results inapplicable. We discuss this further in the next subsection, but for now we use Hölder's inequality to break the dependence. This introduces an additional factor of S which we show how to avoid for small ε in the next section.

$$\begin{aligned}|((\hat{P}_k - P)(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k}| &\le \|(\hat{P}_k - P)(s,a,t)\|_1 \|\tilde{V}_{t+1}^{\pi_k}\|_\infty \\ &\le O\left(\sqrt{\frac{SH^2 \ln \frac{\ln n_{tk}(s,a)}{\delta}}{n_{tk}(s,a)}}\right), \quad (8)\end{aligned}$$

where the second inequality follows from a standard concentration bound. Substituting this into (4) leads to

$$(4) \le O\left(\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sqrt{\frac{SH^2 \ln \frac{\ln n_{tk}(s,a)}{\delta}}{n_{tk}(s,a)}}\right),$$

Now the same argument that we used to obtain (7) shows that the number of episodes when the above display is larger than $\varepsilon/3$ is at most

$$O\left(\frac{S^2AH^4}{\varepsilon^2}\operatorname{polylog}(A,S,H,1/\varepsilon,1/\delta)\right).$$

Combining the arguments we made to control the number of episodes where Eq.$(2,3,4) \ge \varepsilon/3$ explains the second term in the minimum in the statement of the theorem, but not the first, which we do now.

4.2. Improving the Dependence on the State Space

Removing the additional dependence on S requires us to replace the application of Hölder's inequality in (8) with something more refined. A tempting idea is to treat $\hat{P}_k(s,a,t)^\top \tilde{V}_{t+1}^{\pi_k}$ as a weighted sum of independent random variables and apply standard concentration analysis to remove a factor of S in (8). The idea fails, however, because $\tilde{V}_{t+1}^{\pi_k}$ is itself random and depends on *all* the samples obtained so far. This difficulty means that the standard tools of concentration of measure (including martingale methods) are not available to us. Unfortunately this subtle issue invalidates some prior work for which the sample-complexity was stated to be linear in S, where similar terms

were bounded as if $\tilde{V}_{t+1}^{\pi_k}$ were not random (Pazis & Parr, 2016; Pazis et al., 2016). More details are available in Appendix D .

To achieve linear dependence on the state space while addressing the non-independence, one approach is to discard data and introduce phased updating. The model-free Delayed Q-Learning (Strehl et al., 2009) and MoRMax (Szita & Szepesvári, 2010) both take this approach. MoRMax was the first model-based algorithm with a PAC bound linear in S. However, discarding the data and phased updating is unsatisfying from a performance perspective, and indeed, except for their dependence on the size of the state-space, the theoretical guarantees of these two approaches are significantly weaker than Theorem 1 (see also Table 1).

Instead, our approach presented in the following is closer to work by Gheshlaghi Azar et al. (2013) on learning an optimal policy in MDPs with a generative model. The only term which causes difficulty is (4), which we further decompose into

$$\sum_{\substack{t\in[H]\\(s,a)\in L_{tk}}} w_{tk}(s,a)((\hat{P}_k - P)(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k}$$

$$= \sum_{\substack{t\in[H]\\(s,a)\in L_{tk}}} w_{tk}(s,a)((\hat{P}_k - P)(s,a,t))^\top V_{t+1}^{*} \tag{9}$$

$$+ \sum_{\substack{t\in[H]\\(s,a)\in L_{tk}}} w_{tk}(s,a)((\hat{P}_k - P)(s,a,t))^\top (\tilde{V}_{t+1}^{\pi_k} - V_{t+1}^{*}) \tag{10}$$

Since V_{t+1}^* is not random, straightforward concentration analysis is used to bound the first term above by

$$|((\hat{P}_k - P)(s,a,t))^\top V_{t+1}^*| = O\left(\sqrt{\frac{H^2 \ln \frac{\ln n_{tk}(s,a)}{\delta}}{n_{tk}(s,a)}}\right).$$

We can then treat this term in Eq. (9) identically to the term in Eq. (3). Now comes the main point, which is that the remaining term (10) is the product of two terms that each shrink with $\tilde{O}(n_{tk}(s,a)^{-1/2})$ and so the product shrinks like $\tilde{O}(1/n_{tk}(s,a))$. The leading constant hidden by the big-O may have a larger dependence on S than we would like, but this is offset by the faster-than-normal learning rate when ε is reasonably small (and so $n_{tk}(s,a)$ should be large). We now make this argument more precise.

$$|((\hat{P}_k - P)(s,a,t))^\top (\tilde{V}_{t+1}^{\pi_k} - V_{t+1}^{*})|$$

$$\lesssim \sum_{s'\in\mathcal{S}} \sqrt{P(s'|s,a,t)\frac{\ln \frac{\ln n_{tk}(s,a)}{\delta}}{n_{tk}(s,a)}} (\tilde{V}_{t+1}^{\pi_k}(s') - V_{t+1}^{\pi_k}(s'))$$

$$\leq \sqrt{P(s,a,t)^\top (\tilde{V}_{t+1}^{\pi_k} - V_{t+1}^{\pi_k})^2 \frac{S \ln \frac{\ln n_{tk}(s,a)}{\delta}}{n_{tk}(s,a)}}, \tag{11}$$

where we first applied a concentration bound on the empirical mean $\hat{P}_k(s'|s,a,t)$ of Bernoulli random variables and used that $\tilde{V}_{t+1}^{\pi_k}(s') - V_{t+1}^{\pi_k}(s') \geq \tilde{V}_{t+1}^{\pi_k}(s') - V_{t+1}^{*}(s') \geq 0$. The second inequality is an application of Cauchy Schwarz. Notice that the concentration bound was only applied to the empirical error of a transition probability, so the independence issues do not arise. The next step is to show that $P(s,a,t)^\top (\tilde{V}_{t+1}^{\pi_k} - V_{t+1}^{\pi_k})^2 = \tilde{O}(1/n_{tk}(s,a))$. Since $\tilde{V}_{t+1}^{\pi_k} - V_{t+1}^{\pi_k}$ depends on the samples at time $t+1$ to H, we have to relate the number of observations $n_{tk}(s,a)$ of time step t with the number of observations $n_{uk}(s',a')$ of all states and actions at later time steps $u > t$. We leverage the insight that if the conditional probability of observing (s',a') at time step u given $s_t = s, a_t = a$ is sufficiently large, then $n_{uk}(s',a')$ is likely to be of similar magnitude as $n_{tk}(s,a)$ after a few episodes (see Lemma C.4 in the appendix).

Combining $P(s,a,t)^\top (\tilde{V}_{t+1}^{\pi_k} - V_{t+1}^{\pi_k})^2 \lesssim \frac{\ln \frac{\ln n_{tk}(s,a)}{\delta}}{n_{tk}(s,a)}$ with Equation (11) makes the full term in Equation (10) order

$$\sum_{\substack{t\in[H]\\s,a\in L_{tk}}} w_{tk}(s,a) \frac{\mathrm{poly}(S,A,H) \ln \frac{\ln n_{tk}(s,a)}{\delta}}{n_{tk}(s,a)}$$

where $\mathrm{poly}(S,A,H)$ is a polynomial. This quantity decreases faster in $n_{tk}(s,a)$ than the other terms, which decrease like $O(n_{tk}(s,a)^{-1/2})$ and therefore the number of episodes where this term can be larger than $\varepsilon/6$ for small ε is dominated by the number of suboptimal episodes due to other terms. Therefore, for $\varepsilon \leq \frac{1}{AS^2}$ we achieve a total sample complexity of $O\left(\frac{ASH^4}{\varepsilon^2}\mathrm{polylog}(A,S,H,1/\varepsilon,1/\delta)\right)$, which completes the justification for Theorem 1.

4.3. PAC Lower Bound

Theorem 2. *There exist positive constants c, $\delta_0 > 0$, $\varepsilon_0 > 0$ such that for every $\varepsilon \in (0,\varepsilon_0)$, $S \geq 4, A \geq 2$ and for every algorithm A that and $n \leq \frac{cASH^3}{\varepsilon^2}$ there is a fixed-horizon episodic MDP M_{hard} with time-dependent transition probabilities and S states and A actions so that returning an ε-optimal policy after n episodes is at most $1-\delta_0$. That implies that no algorithm can have a PAC guarantee better than $\Omega\left(\frac{ASH^3}{\varepsilon^2}\right)$ for sufficiently small ε.*

Note that this lower bound on the sample complexity of any method in episodic MDPs with time-dependent dynamics applies to the arbitrary but fixed ε PAC bounds and therefore immediately to the stronger δ-uniform PAC bounds. This theorem can be proved in the same way as Theorem 5 by Jiang et al. (2016), which itself is a standard construction involving a careful layering of difficult instances of

the multi-armed bandit problem.[3] For simplicity, we omitted the dependency on the failure probability δ, but using the techniques in the proof of Theorem 26 by Strehl et al. (2009), a lower bound of order $\Omega\left(\frac{ASH^3}{\varepsilon^2}\log(SA/\delta)\right)$ can be obtained. The lower bound shows for small ε the sample complexity of UBEV given in Theorem 1 is optimal except for a factor of H and logarithmic terms.

5. Regret Analysis

The regret of an algorithm is the cumulative difference in expected returns of the optimal policy and the policy followed by the algorithm

$$R(T) = \sum_{k=1}^{T} p_0^\top \left(V_1^* - V_1^{\pi_k}\right).$$

As we remarked in the introduction, an advantage of the new algorithm and δ-uniform-PAC analysis is that it can be used to derive a high probability regret bound. As far as we know, this is the first time a correspondence between PAC and regret guarantees has been formally established. More specifically, Theorem E.1 in the appendix shows that any δ-uniform PAC algorithm also enjoys sub-linear regret with multiplicative constants that depend on its sample-complexity. Applying the theorem to UBEV leads to the following theorem.

Theorem 3. *With probability at least $1-\delta$ the regret of UBEV is bounded jointly for all T by*

$$R(T) = O\left(H^2\sqrt{AST}\,\mathrm{polylog}(T, A, S, H, 1/\delta)\right).$$

It is important to point out that we study the frequentist regret, which is different from the Bayesian regret used by Osband & Van Roy (2016) and others (the Bayesian regret is the integral of $R(T)$ over a class of MDPs). The bound in Theorem 3 is a frequentist bound that holds with high probability for any fixed MDP. In contrast a bound on the Bayesian regret only controls the regret *on average* over all MDPs, so it may be possible that an algorithm minimising the Bayesian regret is sometimes much worse than UBEV (though empirically this is hard to demonstrate).

The best known lower bound for the Bayesian regret in MDPs with-independent dynamics is $\Omega(\sqrt{HAST})$ (see Osband & Van Roy (2016)). This is also a lower bound on the high-probability regret in our setting, though we believe the dependence on H could be increased by exploiting the time-dependence that we permit in our setup.

The uniform nature of our PAC bound and near-optimal dependency of ASH^4 in the $1/\varepsilon^2$ term of the PAC bound is crucial for obtaining the regret bound above. PAC bounds that only hold for fixed ε can only be converted into regret bounds with suboptimal $T^{2/3}$ dependency. In addition the dependency $\sqrt{ASH^4}$ is a direct consequence of the ASH^4/ε^2 dependency of the PAC bound, the $A^2S^3H^4/\varepsilon$ bound contributes to regret in a quantity logarithmic in T.

[3] We here only use $H/2$ timesteps for bandits and the remaining $H/2$ timesteps to accumulate a reward of $O(H)$ for each bandit

6. Empirical Performance

We compared UBEV empirically against other RL methods with PAC or regret guarantees on randomly generated MDPs. We adapted methods that have been developed for a different setting (e.g. γ-discounted MDPs) to episodic MDPs with time-dependent dynamics. While we have not re-derived the theoretical guarantees we expect that potentially guarantees similar to the original setting can be shown. Unlike previous comparisons which run algorithms with empirically tuned parameters that violate the assumptions of the theoretical analyses, we use parameters $\delta = \varepsilon = \frac{1}{10}$ where PAC guarantees hold. Details are in Appendix A .

Figure 1 shows the expected return of the policy followed by each algorithm smoothed over multiple episodes on a representative run for MDPs with $5, 50$ and 200 states. We observe that MedianPAC (Pazis et al., 2016), Delayed Q-Learning (Strehl et al., 2009) and MoRMax (Szita & Szepesvári, 2010) completely fail to improve upon their initial policies within 10M episodes for all but MDPs with very few states and also OIM (Szita & Loerincz, 2008) and UCFH (Dann & Brunskill, 2015) does not find the optimal policy within the episode budget on larger MDPs. UCRL2 (Jaksch et al., 2010) does significantly better but also fails to learn a good policy for 200 states. The very similar MBIE algorithm achieves much better performance since it updates the policy more frequently and more importantly its confidence intervals shrink at rate $n^{-\frac{1}{2}}$ instead UCRL's $\sqrt{\ln(n)/n}$.

Our UBEV algorithm significantly outperforms MBIE and all other PAC methods. UBEV also scales empirically better with the number of states than competitors as predicted by the better PAC guarantee with linear S dependency. For 200 states, UBEV learns a close to optimal policy $10\times$ faster than other PAC algorithms.

Posterior sampling (or PSRL, Osband et al., 2013) converges much faster to an optimal policy, but does not (yet) come with frequentist guarantees. Under weak assumptions it *does* have a Bayesian regret of $\tilde{O}(\sqrt{SAT})$. The question of converting this into a frequentist regret bound is an intriguing and important open problem. The advantage of PSRL is that it does not depend on a tight concentration analysis, which unfortunately is critical for designing practical optimistic algorithms. If we are willing to give up on the frequentist guarantees and scale the confidence bounds

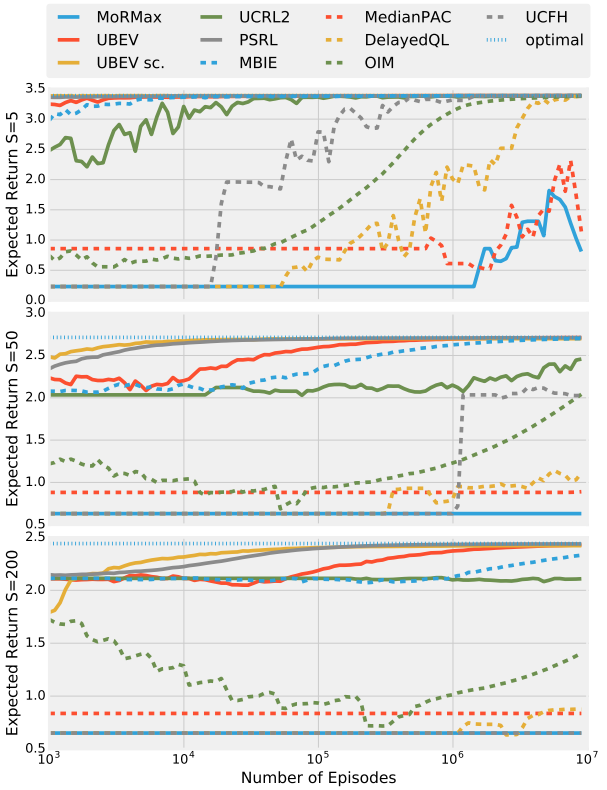

Figure 1: Empirical comparison of PAC algorithms. *UBEV sc.* is a version of UBEV with confidence intervals scaled by $\frac{1}{10}$.

by $1/10$ in Line 7 of Alg. 1, then UBEV matches the performance of PSRL. We have also compared the algorithms on the Chain MDPs used by Osband & Van Roy (2016) and obtained qualitatively similar results.

7. Discussion and Conclusion

We have introduced the UBEV algorithm for episodic MDPs which is simple, efficient and comes with stronger PAC guarantees for small ε than all predecessors. More specifically, other algorithms with linear dependence on the state-space are MoRMax and Delayed Q-Learning, but both of these algorithms have significantly worse dependence on the horizon H. The UCFH algorithm by Dann & Brunskill (2015) improves the H dependence to H^3, but loses the near-linear dependence on S.

We want to emphasize that UBEV is simple and quite practical. For MDPs of reasonable size it can be implemented efficiently. It uses all the data, never stops learning, and updates its policy after every episode. This was achieved by ensuring that all theoretical tricks to make things work only affected the analysis, while in the past many authors modified the algorithm for a simpler analysis.

We have introduced the notion of δ-uniform PAC and shown that UBEV satisfies this stronger sample-complexity guarantee. The uniform nature of our PAC analysis, where the mistake bound holds for all ε simultaneously allowed us to derive a connection between the PAC and regret performance criteria. This means that *without changing our algorithm* we can immediately derive a high probability regret bound of $\tilde{O}(H^2\sqrt{SAT})$, which closes the open question of whether or not the size of the state-space belongs inside the square root (Auer et al., 2010).

Improving the dependence on H The upper and lower PAC bounds differ by a factor of H. An open question is whether or not the techniques in this article can be combined with the observation that exploiting the low variance of near-sparse transitions can be used to shave the last factor of H from the regret (Lattimore & Hutter, 2014; Dann & Brunskill, 2015). The main difficulty is that the arguments in the aforementioned papers rely heavily on controlling the higher-order moments of the value function. But the decomposition argument used in this work requires the value function to be optimistic, but there is no guarantee that this is true for higher-order value functions. Nevertheless, we expect some combination of these ideas will eventually lead to an optimal PAC *and* regret analysis.

Other setups Here we studied the undiscounted episodic MDP setting because of its simplicity. However we expect the same ideas should work for more complicated (and arguably more practical) settings, including the undiscounted infinite-horizon regret setup used by Auer et al. (2010) and others, or the discounted PAC setting of Kakade (2003).

Eliminating lower order terms We are curious to know if the minimum in Theorem 1 can be removed. We do believe that the $\varepsilon S^2 A$ can be replaced by εSA by using a more refined analysis, but removing the dependence entirely seems challenging. Note that other refined analyses also featured such restrictions (Lattimore & Hutter, 2014; Gheshlaghi Azar et al., 2013).

Further improving practicality While the newly proposed algorithm is a step forward for theoretically sound frequentist algorithms, its performance is still generally outperformed by PSRL. One of the great advantages of the pseudo-Bayesian techniques is that they 'auto-tune' to many of the statistical properties of the problem. But in some sequential optimisation problems, such as finite-armed bandits and linear bandits, it is known that carefully designed frequentist approaches can outperform posterior sampling by a significant (sometimes arbitrarily large) margin (Russo & Van Roy, 2014; Lattimore & Szepesári, 2017). The usual approach to improving algorithms based on upper confidence bounds is to shrink the confidence intervals to the absolute minimum. One direction that could be explored is to use informational confidence bounds for multinomial distribution, for example, as used by Filippi

et al. (2010). Though this would need to be modified in some way to include the weighting by the value function.

References

Auer, Peter, Jaksch, Thomas, and Ortner, Ronald. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 99:1563–1600, August 2010. ISSN 1532-4435.

Bellemare, Marc G, Srinivasan, Sriram, Ostrovski, Georg, Schaul, Tom, Saxton, David, Deepmind, Google, and Munos, Rémi. Unifying Count-Based Exploration and Intrinsic Motivation. In *Advances in Neural Information Processing*, 2016.

Boucheron, Stephane, Lugosi, Gabor, and Massart, Pascal. *Concentration Inequalities - A Nonasymptotic Theory of Independence*. Oxford University Press, 2013. ISBN 978-0-19-953525-5.

Dann, Christoph and Brunskill, Emma. Sample Complexity of Episodic Fixed-Horizon Reinforcement Learning. In *Neural Information Processing Systems*, 2015.

Durrett, Rick. *Probability - Theory and Examples*. Cambridge University Press, 4 edition, 2010. ISBN 978-0-521-76539-8.

Filippi, Sarah, Cappé, Olivier, and Garivier, Aurélien. Optimism in reinforcement learning and kullback-leibler divergence. In *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pp. 115–122. IEEE, 2010.

Gheshlaghi Azar, Mohammad, Munos, Rémi, and Kappen, Hilbert J. Minimax PAC bounds on the sample complexity of reinforcement learning with a generative model. *Machine Learning*, 91(3):325–349, 2013.

Jaksch, Thomas, Ortner, Ronald, and Auer, Peter. Near-optimal Regret Bounds for Reinorcement Learning. *Journal of Machine Learning Research*, 11:1563–1600, 2010.

Jiang, Nan, Krishnamurthy, Akshay, Agarwal, Alekh, Langford, John, and Schapire, Robert E. Contextual Decision Processes with Low Bellman Rank are PAC-Learnable, 2016. URL `https://arxiv.org/pdf/1610.09512.pdf`.

Kakade, Sham Machandranath. *On the Sample Complexity of Weak Learning*. PhD thesis, University College London, 2003.

Kearns, Michael and Singh, Satinder. Near-optimal reinforcement learning in polynomial time. *Machine Learning*, 49(2-3):209–232, 2002.

Lattimore, Tor and Hutter, Marcus. Near-optimal PAC bounds for discounted MDPs. In *Theoretical Computer Science*, volume 558, pp. 125–143, 2014. ISBN 9783642341052. doi: 10.1016/j.tcs.2014.09.029.

Lattimore, Tor and Szepesári, Csaba. The end of optimism? An asymptotic analysis of finite-armed linear bandits. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 2017.

Mnih, Volodymyr, Kavukcuoglu, Koray, Silver, David, Rusu, Andrei a, Veness, Joel, Bellemare, Marc G, Graves, Alex, Riedmiller, Martin, Fidjeland, Andreas K, Ostrovski, Georg, Petersen, Stig, Beattie, Charles, Sadik, Amir, Antonoglou, Ioannis, King, Helen, Kumaran, Dharshan, Wierstra, Daan, Legg, Shane, and Hassabis, Demis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015. ISSN 0028-0836. doi: 10.1038/nature14236.

Osband, Ian and Van Roy, Benjamin. Why is Posterior Sampling Better than Optimism for Reinforcement Learning, 2016. URL `http://arxiv.org/abs/1607.00215`.

Osband, Ian, Russo, Daniel, and Van Roy, Benjamin. (More) Efficient Reinforcement Learning via Posterior Sampling. In *Advances in Neural Information Processing*, jun 2013.

Osband, Ian, Van Roy, Benjamin, and Wen, Zheng. Generalization and exploration via randomized value functions. *Proceedings of the 33rd International Conference on Machine Learning*, 2016.

Pazis, Jason and Parr, Ronald. Efficient PAC-optimal Exploration in Concurrent , Continuous State MDPs with Delayed Updates. In *AAAI Conference on Artificial Intelligence*, 2016.

Pazis, Jason, Parr, Ronald E., and How, Jonathan P. Improving PAC Exploration Using the Median Of Means. In *Advances in Neural Information Processing*, 2016.

Russo, Dan and Van Roy, Benjamin. Learning to optimize via information-directed sampling. In *Advances in Neural Information Processing Systems*, pp. 1583–1591, 2014.

Strehl, Alexander L and Littman, Michael L. Exploration via Model-based Interval Estimation. In *International Conference on Machine Learning*, number Icml, 2004.

Strehl, Alexander L, Li, Lihong, and Littman, Michael L. Reinforcement Learning in Finite MDPs : PAC Analysis. *Journal of Machine Learning Research*, 10:2413–2444, 2009.

Szita, Istvàn and Loerincz, Andras. The Many Faces of Optimism : a Unifying Approach. In *International Conference on Machine Learning*, 2008.

Szita, Istvàn and Szepesvári, Csaba. Model-based reinforcement learning with nearly tight exploration complexity bounds. In *International Conference on Machine Learning*, 2010.

Appendices

A. Experimental Details

We generated the MDPs with $S = 5, 50, 200$ states, $A = 3$ actions and $H = 10$ timesteps as follows: The transition probabilities $P(s, a, t)$ were sampled independently from Dirichlet $\left(\frac{1}{10}, \dots \frac{1}{10}\right)$ and the rewards were all deterministic with their value $r(s, a, t)$ set to 0 with probability 85% and set uniformly at random in $[0, 1]$ otherwise. This construction results in MDPs that have concentrated but non-deterministic transition probabilities and sparse rewards.

Since some alorithms have been proposed assuming the rewards $r(s, a, t)$ are known and we aim for a fair comparison, we assumed for all algorithms that the immediate rewards $r(s, a, t)$ are known and adapted the algorithms accordingly. For example, in UBEV, the $\min\left\{1, \frac{l(s,a,t)}{\max\{1,n(s,a,t)\}} + \phi\right\}$ term was replaced by the true known rewards $r(s, a, t)$ and the δ parameter in ϕ was scaled by $9/7$ accordingly since the concentration result for immediate rewards is not necessary in this case. We used $\delta = \frac{1}{10}$ for all algorithms and $\varepsilon = \frac{1}{10}$ if they require to know ε beforehand.

We adapted MoRMax, UCRL2, UCFH, MBIE, PSRL MeidanPAC, Delayed Q-Learning and OIM to the episodic MDP setting with time-dependent transition dynamics by using allowing them to learn time-dependent dynamics and use finite-horizon planning. We did adapt the confidence intervals and but did not rederive the constants for each algorithm. When in doubt we opted for smaller constants typically resulting better performance of the competitors. We also reduced the number of episodes used in the delays by a factor of $\frac{1}{1000}$ for MoRMax and Delayed Q-Learning and by 10^{-6} for UCFH because they would otherwise not have performed a single policy update even for $S = 5$ within the 10 million episodes we considered. This scaling violates their theoretical guarantees but at least shows that the methods work in principle.

The performance reported in Figure 1 are the expected return of the current policy of each algorithm averaged over 1000 episodes. The figure shows a single run of the same randonly generated MDP but the results are representative. We reran this experiments with different random seeds and consistently obtained qualitatively similar results.

Source code for all experiments including concise but efficient implementations of algorithms is available at `https://github.com/chrodan/FiniteEpisodicRL.jl`.

B. Planning Problem of UBEV

Lemma B.1 (Planning Problem)**.** *The policy update in Lines 4–10 of Algorithm 1 finds an optimal solution to the optimization problem*

$$
\begin{aligned}
&\max_{P', V', \pi', r'} \mathbb{E}_{s\sim p_0}[V'_1(s)] \\
&\forall s \in \mathcal{S}, a \in \mathcal{A}, t \in [H]: \\
&\quad V'_{H+1} = 0, \qquad P'(s,a,t) \in \Delta_S, \qquad r'(s,a,t) \in [0,1] \\
&\quad V'_t(s) = r'(s, \pi'(s,t), t) + \mathbb{E}_{s'\sim P'(s,\pi'(s,t),t)}[V'_{t+1}] \\
&\quad |(P'(s,a,t) - \hat{P}_k(s,a,t))^\top V'_{t+1}| \le \phi(s,a,t) \operatorname{rng} V'_{t+1} \\
&\quad |r'(s,a,t) - \hat{r}_k(s,a,t)| \le \phi(s,a,t)
\end{aligned}
$$

where $\phi(s,a,t) = \sqrt{\frac{2\operatorname{llnp}(n(s,a,t)) + \ln(18SAH/\delta)}{n(s,a,t)}}$ *is a confidence bound and* $\hat{P}_k(s'|s,a,t) = m(s',s,a,t)/n(s,a,t)$ *are the empirical transition probabilities and* $\hat{r}_k(s,a,t) = l(s,a,t)/n(s,a,t)$ *the empirical average rewards.*

Proof. Since $\tilde{V}_{H+1}(\cdot)$ is initialized with 0 and never changed, we immediately get that it is an optimal value for $V'_{H+1}(\cdot)$ which is constrained to be 0. Consider now a single time step t and assume V'_{t+1} are fixed to the optimal values $\tilde{V}_{t+1}$. Plugging in the computation of $EV(a)$ into the computation of $\tilde{V}_t(s)$, we get

$$
\begin{aligned}
\tilde{V}_t(s) = \max_a EV(a) = \max_{a\in\mathcal{A}} \Bigg[&\min\{1, \hat{r}(s,a,t) + \phi(s,a,t)\} \\
&+ \min\left\{\max \tilde{V}_{t+1}, \mathbb{I}\{n(s,a,t) > 0\}(\hat{P}(s,a,t)^\top \tilde{V}_{t+1}) + \phi(s,a,t) \operatorname{rng} \tilde{V}_{t+1}\right\}\Bigg]
\end{aligned}
$$

using the convention that $\hat{r}(s,a,t) = 0$. Assuming that $V'_{t+1} = \tilde{V}_{t+1}$, and that our goal for now is to maximize $\tilde{V}_t(s,)$, this can be rewritten as

$$\begin{aligned}\max_{P'(s,a,t),r'(s,a,t)} \tilde{V}_t(s) = \max_{P'(s,a,t),r'(s,a,t),\pi'(s,t)} &\left[r'(s,\pi'(s,t),t) + P'(s,\pi'(s,t),t)^\top \tilde{V}_{t+1}\right] \\ \text{s.t.} \qquad &\forall a \in \mathcal{A} : r'(s,a,t) \in [0,1], \qquad P'(s,a,t) \in \Delta_S \\ &|(P'(s,a,t) - \hat{P}_k(s,a,t))^\top V'_{t+1}| \leq \phi(s,a,t) \operatorname{rng} V'_{t+1} \\ &|r'(s,a,t) - \hat{r}_k(s,a,t)| \leq \phi(s,a,t)\end{aligned}$$

since in this problem either $P'(s,\pi'(s,t),t)^\top \tilde{V}_{t+1} = \hat{P}(s,\pi'(s,t),t)^\top \tilde{V}_{t+1} + \phi(s,a,t) \operatorname{rng} V'_{t+1}$ if that does not violate $P'(s,\pi'(s,t),t)^\top \tilde{V}_{t+1} \leq \max \tilde{V}_{t+1}$ and otherwise $P'(s',s,\pi'(s,t),t) = 1$ for one state s' with $\tilde{V}_{t+1}(s') = \max \tilde{V}_{t+1}$. Similarly, either $r'(s,\pi'(s,t),t) = \hat{r}(s,\pi'(s,t),t) + \phi(s,\pi'(s,t),t)$ if that does not violate $r'(s,\pi'(s,t),t) \leq 1$ or $r'(s,\pi'(s,t),t) = 1$ otherwise. Using induction for $t = H, H-1 \dots 1$, we see that UBEV computes an optimal solution to

$$\begin{aligned}&\max_{P',V',\pi',r'} V'_1(\tilde{s}) \\ &\forall s \in \mathcal{S}, a \in \mathcal{A}, t \in [H] : \\ &\quad V'_{H+1} = 0, \qquad P'(s,a,t) \in \Delta_S, \qquad r'(s,a,t) \in [0,1] \\ &\quad V'_t(s) = r'(s,\pi'(s,t),t) + \mathbb{E}_{s' \sim P'(s,\pi'(s,t),t)}[V'_{t+1}] \\ &\quad |(P'(s,a,t) - \hat{P}_k(s,a,t))^\top V'_{t+1}| \leq \phi(s,a,t) \operatorname{rng} V'_{t+1} \\ &\quad |r'(s,a,t) - \hat{r}_k(s,a,t)| \leq \phi(s,a,t)\end{aligned}$$

for any fixed $\tilde{s}$. The intersection of all optimal solutions to this problem for all $\tilde{s} \in \mathcal{S}$ are also an optimal solution to

$$\begin{aligned}&\max_{P',V',\pi',r'} p_0^\top V'_1 \\ &\forall s \in \mathcal{S}, a \in \mathcal{A}, t \in [H] : \\ &\quad V'_{H+1} = 0, \qquad P'(s,a,t) \in \Delta_S, \qquad r'(s,a,t) \in [0,1] \\ &\quad V'_t(s) = r'(s,\pi'(s,t),t) + \mathbb{E}_{s' \sim P'(s,\pi'(s,t),t)}[V'_{t+1}] \\ &\quad |(P'(s,a,t) - \hat{P}_k(s,a,t))^\top V'_{t+1}| \leq \phi(s,a,t) \operatorname{rng} V'_{t+1} \\ &\quad |r'(s,a,t) - \hat{r}_k(s,a,t)| \leq \phi(s,a,t).\end{aligned}$$

Hence, UBEV computes an optimal solution to this problem. □

C. Details of PAC Analysis

In the analysis, we denote the value of $n(\cdot,t)$ after the planning in iteration k as $n_{tk}(\cdot)$. We further denote by $P(s'|s,a,t)$ the probability of sampling state s' as s_{t+1} when $s_t = s, a_t = a$. With slight abuse of notation, $P(s,a,t) \in [0,1]^S$ denotes the probability vector of $P(\cdot|s,a,t)$. We further use $\tilde{P}_k(s'|s,a,t)$ as conditional probability of $s_{t+1} = s'$ given $s_t = s, a_t = a$ but in the optimistic MDP $\tilde{M}$ computed in the optimistic planning steps in iteration k. We also use the following definitions:

$$\begin{aligned}w_{\min} &= w'_{\min} = \frac{\varepsilon c_\varepsilon}{H^2 S} \\ c_\varepsilon &= \frac{1}{3} \\ L_{tk} &= \{(s,a) \in \mathcal{S} \times \mathcal{A} : w_{tk}(s,a) \geq w_{\min}\} \\ \operatorname{llnp}(x) &= \ln(\ln(\max\{x,e\})) \\ \delta' &= \frac{\delta}{9}\end{aligned}$$

In the following, we provide the formal proof for Theorem 1 and then present all necessary lemmas:

C.1. Proof of Theorem 1

Proof of Theorem 1. Corollary C.5 ensures that the failure event has probability at most δ. Outside the failure event Lemma C.2 ensures that all but at most $\frac{48A^2S^3H^4}{\varepsilon}\operatorname{polylog}(A,S,H,1/\varepsilon,1/\delta)$ episodes are friendly. Finally, Lemma C.8 shows that all friendly episodes except at most $\left(\frac{9216}{\varepsilon}+417S\right)\frac{ASH^4}{\varepsilon}\operatorname{polylog}(A,S,H,1/\varepsilon,1/\delta)$ are ε-optimal. The second bound follows from replacing AS^2 by $1/\varepsilon$ in the second term. Furthermore, outside the failure event Lemma C.2 ensures that all but at most $\frac{6AS^2H^3}{\varepsilon}\operatorname{polylog}(A,S,H,1/\varepsilon,1/\delta)$ episodes are nice. Finally, Lemma C.7 shows that all nice episodes except at most $(4+S)\,576\frac{ASH^4}{\varepsilon}\operatorname{polylog}(A,S,H,1/\varepsilon,1/\delta)$ are ε-optimal.

□

C.2. Failure Events and Their Probabilities

In this section, we define a failure event F in which we cannot guarantee the performance of UBEV. We then show that this event F only occurs with low probability. All our arguments are based on general uniform concentration of measure statements that we prove in Section F. In the following we argue how the apply in our setting and finally combine all concentration results to get $\mathbb{P}(F)\le\delta$. The failure event is defined as

$$F=\bigcup_k\left[F_k^N\cup F_k^{CN}\cup F_k^P\cup F_k^V\cup F_k^{L1}\cup F_k^R\right]$$

where

$$\begin{aligned}
F_k^N &= \left\{\exists s,a,t:\ n_{tk}(s,a)<\frac{1}{2}\sum_{i<k}w_{ti}(s,a)-\ln\frac{SAH}{\delta'}\right\}\\
F_k^{CN} &= \left\{\exists s,a,s',a',u<t:\ n_{tk}(s,a)<\frac{1}{2}n_{uk}(s',a')\sum_{i<k}w_{ui}^t(s,a|s',a')-\ln\left(\frac{S^2A^2H^2}{\delta'}\right)\right\}\\
F_k^V &= \left\{\exists s,a,t:\ |(\hat P_k(s,a,t)-P(s,a,t))^\top V_{t+1}^*|\ge\sqrt{\frac{\operatorname{rng}(\tilde V_{t+1}^*)^2}{n_{tk}(s,a)}\left(2\operatorname{llnp}(n_{tk}(s,a))+\ln\frac{3SAH}{\delta'}\right)}\right\}\\
F_k^P &= \left\{\exists s,s',a,t:\ |\hat P_k(s'|s,a,t)-P(s'|s,a,t)|\ge\sqrt{\frac{2P(s'|s,a,t)}{n_{tk}(s,a)}\left(2\operatorname{llnp}(n_{tk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)}\right.\\
&\qquad\left.+\frac{1}{n_{tk}(s,a)}\left(2\operatorname{llnp}(n_{tk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)\right\}\\
F_k^{L1} &= \left\{\exists s,a,t:\ \|\hat P_k(s,a,t)-P(s,a,t)\|_1\ge\sqrt{\frac{4}{n_{tk}(s,a)}\left(2\operatorname{llnp}(n_{tk}(s,a))+\ln\frac{3SAH(2^S-2)}{\delta'}\right)}\right\}\\
F_k^R &= \left\{\exists s,a,t:\ |\hat r_k(s,a,t)-r(s,a,t)|\ge\sqrt{\frac{1}{n_{tk}(s,a)}\left(2\operatorname{llnp}(n_{tk}(s,a))+\ln\frac{3SAH}{\delta'}\right)}\right\}.
\end{aligned}$$

We now bound the probability of each type of failure event individually:

Corollary C.1. *For any $\delta'>0$, it holds that $\mathbb{P}\left(\bigcup_{k=1}^\infty F_k^V\right)\le 2\delta'$ and $\mathbb{P}\left(\bigcup_{k=1}^\infty F_k^R\right)\le 2\delta'$*

Proof. Consider a fix $s\in\mathcal{S},a\in\mathcal{A},t\in[H]$ and denote $\mathcal{F}_k$ the sigma-field induced by the first $k-1$ episodes and the k-th episode up to s_t and a_t but not s_{t+1}. Define τ_i to be the index of the episode where (s,a) was observed at time t the ith time. Note that τ_i are stopping times with respect to $\mathcal{F}_i$. Define now the filtration $\mathcal{G}_i=\mathcal{F}_{\tau_i}=\{A\in\mathcal{F}_\infty:\ A\cap\{\tau_i\le t\}\in\mathcal{F}_t\ \forall\, t\ge 0\}$ and $X_k=(V_{t+1}^*(s_k')-P(s,a,t)^\top V_{t+1}^*)\mathbb{I}\{\tau_k<\infty\}$ where s_i' is the value of s_{t+1} in episode τ_i (or arbitrary, if $\tau_i=\infty$).

By the Markov property of the MDP, we have that X_i is a martingale difference sequence with respect to the filtration $\mathcal{G}_i$. Further, since $\mathbb{E}[X_i|\mathcal{G}_{i-1}]=0$ and $|X_i|\in[0,\operatorname{rng}(V_{t+1}^*)]$, X_i conditionally $\operatorname{rng}(V_{t+1}^*)/2$-subgaussian due to Hoeffding's Lemma, i.e., satisfies $\mathbb{E}[\exp(\lambda X_i)|\mathcal{G}_{i-1}]\le\exp(\lambda^2\operatorname{rng}(V_{t+1}^*)^2/2)$.

We can therefore apply Lemma F.1 and conclude that

$$\mathbb{P}\left(\exists k : |(\hat{P}_k(s,a,t) - P(s,a,t))^\top V^*_{t+1}| \geq \sqrt{\frac{\mathrm{rng}(V^*_{t+1})^2}{n_{tk}(s,a)}\left(2\,\mathrm{llnp}(n_{tk}(s,a)) + \ln\frac{3}{\delta'}\right)}\right) \leq 2\delta'\,.$$

Analogously

$$\mathbb{P}\left(\exists k : |\hat{r}_k(s,a,t) - r(s,a,t)| \geq \sqrt{\frac{1}{n_{tk}(s,a)}\left(2\,\mathrm{llnp}(n_{tk}(s,a)) + \ln\frac{3}{\delta'}\right)}\right) \leq 2\delta'\,.$$

Applying the union bound over all $s \in \mathcal{S}, a \in \mathcal{A}$ and $t \in [H]$, we obtain the desired statement for F^V. In complete analogy using the same filtration, we can show the statement for F^R. □

Corollary C.2. *For any $\delta' > 0$, it holds that $\mathbb{P}\left(\bigcup_{k=1}^\infty F_k^P\right) \leq 2\delta'$.*

Proof. Consider first a fix $s', s \in \mathcal{S}, t \in [H]$ and $a \in \mathcal{A}$. Let K denote the number of times the triple s, a, t was encountered in total during the run of the algorithm. Define the random sequence X_i as follows. For $i \leq K$, let X_i be the indicator of whether s' was the next state when s, a, t was encountered the ith time and for $i > K$, let $X_i \sim \mathrm{Bernoulli}(P(s'|s,a,t))$ be drawn i.i.d. By construction this is a sequence of i.i.d. Bernoulli random variables with mean $P(s'|s,a,t)$. Further the event

$$\bigcup_k \Bigg\{ \left|\hat{P}_k(s'|s,a,t) - P(s'|s,a,t)\right| \geq \sqrt{\frac{2P(s'|s,a,t)}{n_{tk}(s,a)}\left(2\,\mathrm{llnp}(n(s,a,t)) + \ln\frac{3S^2AH}{\delta'}\right)} + \frac{1}{n_{tk}(s,a)}\left(2\,\mathrm{llnp}(n_{tk}(s,a)) + \ln\frac{3S^2AH}{\delta'}\right) \Bigg\}$$

is contained in the event

$$\bigcup_i \left\{ |\hat{\mu}_i - \mu| \geq \sqrt{\frac{2\mu}{i}\left(2\,\mathrm{llnp}(i) + \ln\frac{3}{\delta'}\right)} + \frac{1}{i}\left(2\,\mathrm{llnp}(i) + \ln\frac{3S^2AH}{\delta'}\right) \right\}$$

whose probability can be bounded by $2\delta'/S^2/A/H$ using Lemma F.2. The statement now follows by applying the union bound. □

Corollary C.3. *For any $\delta' > 0$, it holds that $\mathbb{P}\left(\bigcup_{k=1}^\infty F_k^{L1}\right) \leq \delta'$*

Proof. Using the same argument as in the proof of Corollary C.2 the statement follows from Lemma F.3. □

Corollary C.4. *It holds that*

$$\mathbb{P}\left(\bigcup_k F_k^N\right) \leq \delta' \quad \textit{and} \quad \mathbb{P}\left(\bigcup_k F_k^{CN}\right) \leq \delta'.$$

Proof. Consider a fix $s \in \mathcal{S}, a \in \mathcal{A}, t \in [H]$. We define $\mathcal{F}_k$ to be the sigma-field induced by the first $k-1$ episodes and X_k as the indicator whether s, a, t was observed in episode k. The probability $w_{tk}(s,a)$ pf whether $X_k = 1$ is F_k measurable and hence we can apply Lemma F.4 with $W = \ln\frac{SAH}{\delta'}$ and obtain that $\mathbb{P}\left(\bigcup_k F_k^N\right) \leq \delta'$ after applying the union bound.

For the second statement, consider again a fix $s, s' \in \mathcal{S}, a, a' \in \mathcal{A}, u, t \in [H]$ with $u < t$ and denote by $\mathcal{F}_k$ the sigma-field induced by the first $k-1$ episodes and the k-th episode up to s_u and a_u but not s_{u+1}. Define τ_i to be the index of the episode where (s', a') was observed at time u the ith time. Note that τ_i are stopping times with respect to $\mathcal{F}_i$. Define now the filtration $\mathcal{G}_i = \mathcal{F}_{\tau_i} = \{A \in \mathcal{F}_\infty : A \cap \{\tau_i \leq k\} \in \mathcal{F}_k \ \forall\, k \geq 0\}$ and X_i to be the indicator whether s, a, t and s', a', u was observed in episode τ_i. If $\tau_i = \infty$, we set $X_i = 0$. Note that the probablity $w^t_{ui}(s,a|s',a')\mathbb{I}\{\tau_i < \infty\}$ of $X_i = 1$ is $\mathcal{G}_i$-measureable.

By the Markov property of the MDP, we have that X_i is a martingale difference sequence with respect to the filtration $\mathcal{G}_i$. We can therefore apply Lemma F.4 with $W = \ln \frac{S^2A^2H^2}{\delta'}$ and using the union bound over all s, a, s', a', u, t, we get $\mathbb{P}\left(\bigcup_k F_k^{CN}\right) \leq \delta'$. □

Corollary C.5. *The total failure probability of the algorithm is bounded by* $\mathbb{P}(F) \leq 9\delta' = \delta$.

Proof. Statement follows directly from Corollary C.1, Corollary C.2, Corollary C.3, Corollary C.4 and the union bound. □

C.3. Nice and Friendly Episodes

We now define the notion of *nice* and the stronger *friendly* episodes. In nice episodes, all states either have low probability of occuring or the sum of probability of occuring in the previous episodes is large enough so that outside the failure event we can guarantee that

$$n_{tk}(s,a) \geq \frac{1}{4}\sum_{i<k} w_{ti}(s,a).$$

This allows us to then bound the number of nice episodes terms of the form

$$\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sqrt{\frac{\operatorname{llnp}(n_{tk}(s,a))+D}{n_{tk}(s,a)}}$$

can exceed a chosen threshold (see Lemma C.3 below). In the next section, we will bound the optimality gap of an episode by terms of such form and use the results derived here to bound the number of nice episodes where the algorithm can follow a ε-suboptimal policy. Together with a bound on the number of non-nice episodes, we obtain the sample complexity of UBEV shown in Theorem 1.

Similarly, we use a more refined analysis of the optimality gap of friendly episodes together with Lemma C.4 below to obtain the tighter sample complexity linear-polylog in S.

Definition 1 (Nice and Friendly Episodes)**.** *An episode k is* nice *if and only if for all $s \in \mathcal{S}$, $a \in \mathcal{A}$ and $t \in [H]$ the following two conditions hold:*

$$w_{tk}(s,a) \leq w_{\min} \quad \vee \quad \frac{1}{4}\sum_{i<k} w_{ti}(s,a) \geq \ln\frac{SAH}{\delta'}$$

An episode k is friendly *if and only if it is nice and for all $s, s' \in \mathcal{S}$, $a, a' \in \mathcal{A}$ and $u, t \in [H]$ with $u < t$ the following two conditions hold:*

$$w^t_{uk}(s,a|s',a') \leq w'_{\min} \quad \vee \quad \frac{1}{4}\sum_{i<k} w^t_{ui}(s,a|s',a') \geq \ln\frac{S^2A^2H^2}{\delta'}.$$

We denote the set of all nice episodes by $N \subseteq \mathbb{N}$ and the set of all friendly episodes by $K \subseteq N$.

Lemma C.1 (Properties of nice and friendly episodes)**.** *If an episode k is nice, i.e., $k \in N$, then on F^c (outside the failure event) for all $s \in \mathcal{S}$, $a \in \mathcal{A}$ and $t \in [H]$ with $u < t$ the following statement holds:*

$$w_{tk}(s,a) \leq w_{\min} \quad \vee \quad n_{tk}(s,a) \geq \frac{1}{4}\sum_{i<k} w_{ti}(s,a).$$

If an episode k is friendly, i.e., $k \in K$, then on F^c (outside the failure event) for all $s, s' \in \mathcal{S}$, $a, a' \in \mathcal{A}$ and $u, t \in [H]$ with $u < t$ the above statement holds as well as

$$w^t_{uk}(s,a|s',a') \leq w'_{\min} \quad \vee \quad n_{tk}(s,a) \geq \frac{1}{4}n_{uk}(s',a')\sum_{i<k} w^t_{ui}(s,a|s',a').$$

Proof. Since we consider the event $F_k^{N^c}$, it holds for all s, a, t triples with $w_{tk}(s,a) > w_{\min}$

$$n_{tk}(s,a) \geq \frac{1}{2}\sum_{i<k} w_{ti}(s,a) - \ln\frac{SAH}{\delta'} \geq \frac{1}{4}\sum_{i<k} w_{ti}(s,a)$$

for $k \in N$ Further, since we only consider the event $F_k^{CN^c}$,we have for all $s, s' \in \mathcal{S}$, $a, a' \in \mathcal{A}$, $u, t \in [H]$ with $u < t$ and $w_{uk}^t(s,a|s',a') > w_{\min}$

$$n_{tk}(s,a) \geq \frac{1}{2}n_{uk}(s',a')\sum_{i<k} w_{ui}^t(s,a|s',a') - \ln\frac{S^2A^2H^2}{\delta'}$$

for $k \in E$. If $n_{uk}(s',a') = 0$ then $n_{tk}(s,a) \geq 0 = \frac{1}{4}n_{uk}(s',a')\sum_{i<k} w_{ui}^t(s,a|s',a')$ holds trivially. Otherwise $n_{uk}(s',a') \geq 1$ and therefore

$$\begin{aligned}
n_{tk}(s,a) \geq& \frac{1}{2}n_{uk}(s',a')\sum_{i<k} w_{ui}^t(s,a|s',a') - \ln\frac{S^2A^2H^2}{\delta'} \\
\geq& \frac{1}{2}n_{uk}(s',a')\sum_{i<k} w_{ui}^t(s,a|s',a') - \frac{1}{4}\sum_{i<k} w_{ui}^t(s,a|s',a') \\
\geq& \frac{1}{4}n_{uk}(s',a')\sum_{i<k} w_{ui}^t(s,a|s',a')
\end{aligned}$$

□

Lemma C.2 (Number of non-nice and non-friendly episodes)**.** *On the good event F^c, the number of episodes that are not friendly is at most*

$$48\frac{S^3A^2H^4}{\varepsilon}\ln\frac{S^2A^2H^2}{\delta'}$$

and the number episodes that are not nice is at most

$$\frac{6S^2AH^3}{\varepsilon}\ln\frac{SAH}{\delta'}.$$

Proof. If an episode k is not nice, then there is s, a, t with $w_{tk}(s,a) > w_{\min}$ and $\sum_{i<k} w_{ti}(s,a) < 4\ln\frac{SAH}{\delta'}$. Since the sum on the left-hand side of this inequality increases by at least $w_{\min}$ this happens while the right hand side stays constant, this situation can occur at most

$$\frac{4SAH}{w_{\min}}\ln\frac{SAH}{\delta'} = \frac{24S^2AH^3}{\varepsilon}\ln\frac{SAH}{\delta'}$$

times in total. If an episode k is not friendly, it is either not nice or there is s, a, t and s', a', u with $u < t$ and $w_{uk}^t(s',a'|s,a) > w'_{\min}$ and $\sum_{i<k} w_{ui}^t(s,a|s',a') < 4\ln\frac{S^2A^2H^2}{\delta'}$. Since the sum on the left-hand side of this inequality increases by at least $w'_{\min}$ each time this happens while the right hand side stays constant, this can happen at most $\frac{4S^2A^2H^2}{w'_{\min}}\ln\frac{S^2A^2H^2}{\delta'}$ times in total. Therefore, there can only be at most

$$\begin{aligned}
&\frac{4SAH}{w_{\min}}\ln\frac{SAH}{\delta'} + \frac{4S^2A^2H^2}{w'_{\min}}\ln\frac{S^2A^2H^2}{\delta'} \\
=&\frac{4S^2AH^3}{c_\varepsilon\varepsilon}\ln\frac{SAH}{\delta'} + \frac{4S^3A^2H^4}{c_\varepsilon\varepsilon}\ln\frac{S^2A^2H^2}{\delta'} \leq \frac{48S^3A^2H^4}{\varepsilon^2}\ln\frac{S^2A^2H^2}{\delta'}
\end{aligned}$$

non-friendly episodes. □

Lemma C.3 (Main Rate Lemma)**.** *Let $r \geq 1$ fix and $C > 0$ which can depend polynomially on the relevant quantities and $\varepsilon' > 0$ and let $D \geq 1$ which can depend poly-logarithmically on the relevant quantities. Then*

$$\sum_t \sum_{s,a \in L_{tk}} w_{tk}(s,a) \left(\frac{C(\mathrm{llnp}(n_{tk}(s,a)) + D)}{n_{tk}(s,a)} \right)^{1/r} \leq \varepsilon'$$

on all but at most

$$\frac{8CASH^r}{\varepsilon'^r} \operatorname{polylog}(S, A, H, \delta^{-1}, \varepsilon'^{-1}).$$

nice episodes.

Proof. Define

$$\begin{aligned}
\Delta_k =& \sum_t \sum_{s,a \in L_{tk}} w_{tk}(s,a) \left(\frac{C(\mathrm{llnp}(n_{tk}(s,a)) + D)}{n_{tk}(s,a)} \right)^{1/r} \\
=& \sum_t \sum_{s,a \in L_{tk}} w_{tk}(s,a)^{1-\frac{1}{r}} \left(w_{tk}(s,a) \frac{C(\mathrm{llnp}(n_{tk}(s,a)) + D)}{n_{tk}(s,a)} \right)^{1/r}.
\end{aligned}$$

We first bound using Hölder's inequality

$$\Delta_k \leq \left(\sum_t \sum_{s,a \in L_{tk}} \frac{CH^{r-1} w_{tk}(s,a)(\mathrm{llnp}(n_{tk}(s,a)) + D)}{n_{tk}(s,a)} \right)^{\frac{1}{r}}.$$

Using the property in Lemma C.1 of nice episodes as well as the fact that $w_{tk}(s,a) \leq 1$ and $\sum_{i<k} w_{ti}(s,a) \geq 4 \ln \frac{SAH}{\delta'} \geq 4 \ln(2) \geq 2$, we bound

$$n_{tk}(s,a) \geq \frac{1}{4} \sum_{i<k} w_{ti}(s,a) \geq \frac{1}{8} \sum_{i \leq k} w_{ti}(s,a).$$

The function $\frac{\mathrm{llnp}(x)+D}{x}$ is monotonically decreasing in $x \geq 0$ since $D \geq 1$ (see Lemma C.6). This allows us to bound

$$\begin{aligned}
\Delta_k^r \leq& \sum_t \sum_{s,a \in L_{tk}} \frac{CH^{r-1} w_{tk}(s,a)(\mathrm{llnp}(n_{tk}(s,a)) + D)}{n_{tk}(s,a)} \\
\leq& 8CH^{r-1} \sum_t \sum_{s,a \in L_{tk}} \frac{w_{tk}(s,a) \left(\mathrm{llnp}\left(\frac{1}{8} \sum_{i \leq k} w_{ti}(s,a) \right) + D \right)}{\sum_{i \leq k} w_{ti}(s,a)} \\
\leq& 8CH^{r-1} \sum_t \sum_{s,a \in L_{tk}} \frac{w_{tk}(s,a) \left(\mathrm{llnp}\left(\sum_{i \leq k} w_{ti}(s,a) \right) + D \right)}{\sum_{i \leq k} w_{ti}(s,a)}.
\end{aligned}$$

Assume now $\Delta_k > \varepsilon'$. In this case the right-hand side of the inequality above is also larger than ε'^r and there is at least one (s,a,t) with $w_{tk}(s,a) > w_{\min}$ and

$$\begin{aligned}
&\frac{8CSAH^r \left(\mathrm{llnp}\left(\sum_{i \leq k} w_{ti}(s,a) \right) + D \right)}{\sum_{i \leq k} w_{ti}(s,a)} > \varepsilon'^r \\
&\Leftrightarrow \frac{\mathrm{llnp}\left(\sum_{i \leq k} w_{ti}(s,a) \right) + D}{\sum_{i \leq k} w_{ti}(s,a)} > \frac{\varepsilon'^r}{8CSAH^r}.
\end{aligned}$$

Let us denote $C' = \frac{8CASH^r}{\varepsilon'^r}$. Since $\frac{\text{llnp}(x)+D}{x}$ is monotonically decreasing and $x = C'^2 + 3C'D$ satisfies $\frac{\text{llnp}(x)+D}{x} \leq \frac{\sqrt{x}+D}{x} \leq \frac{1}{C'}$, we know that if $\sum_{i\leq k} w_{ti}(s,a) \geq C'^2 + 3C'D$ then the above condition cannot be satisfied for s, a, t. Since each time the condition is satisfied, it holds that $w_{tk}(s,a) > w_{\min}$ and so $\sum_{i\leq k} w_{ti}(s,a)$ increases by at least $w_{\min}$, it can happen at most

$$m \leq \frac{ASH(C'^2 + 3C'D)}{w_{\min}}$$

times that $\Delta_k > \varepsilon'$. Define $K = \{k : \Delta_k > \varepsilon'\} \cap N$ and we know that $|K| \leq m$. Now we consider the sum

$$\begin{aligned}\sum_{k\in K} \Delta_k^r \leq& \sum_{k\in K} 8CH^{r-1} \sum_t \sum_{s,a\in L_{tk}} \frac{w_{tk}(s,a)\left(\text{llnp}\left(\sum_{i\leq k} w_{ti}(s,a)\right) + D\right)}{\sum_{i\leq k} w_{ti}(s,a)} \\ \leq& 8CH^{r-1}\left(\text{llnp}\left(C'^2 + 3C'D\right) + D\right) \sum_t \sum_{s,a\in L_{tk}} \sum_{k\in K} \frac{w_{tk}(s,a)}{\sum_{i\leq k} w_{ti}(s,a)\mathbb{I}\{w_{ti}(s,a) \geq w_{\min}\}}\end{aligned}$$

For every (s,a,t), we consider the sequence of $w_{ti}(s,a) \in [w_{\min}, 1]$ with $i \in I = \{i \in \mathbb{N} : w_{ti}(s,a) \geq w_{\min}\}$ and apply Lemma C.5. This yields that

$$\sum_{k\in K} \frac{w_{tk}(s,a)}{\sum_{i\leq k} w_{ti}(s,a)\mathbb{I}\{w_{ti}(s,a) \geq w_{\min}\}} \leq 1 + \ln(m/w_{\min}) = \ln\left(\frac{me}{w_{\min}}\right)$$

and hence

$$\sum_{k\in K} \Delta_k^r \leq 8CASH^r \ln\left(\frac{me}{w_{\min}}\right)\left(\text{llnp}\left(C'^2 + 3C'D\right) + D\right)$$

Since each element in K has to contribute at least ε'^r to this bound, we can conclude that

$$\sum_{k\in N} \mathbb{I}\{\Delta_k \geq \varepsilon'\} \leq \sum_{k\in K} \mathbb{I}\{\Delta_k \geq \varepsilon'\} \leq |K| \leq \frac{8CASH^r}{\varepsilon'^r} \ln\left(\frac{me}{w_{\min}}\right)\left(\text{llnp}\left(C'^2 + 3C'D\right) + D\right).$$

Since $\ln\left(\frac{me}{w_{\min}}\right)\left(\text{llnp}\left(C'^2 + 3C'D\right) + D\right)$ is $\text{polylog}(S, A, H, \delta^{-1}, \varepsilon'^{-1})$, the proof is complete. □

Lemma C.4 (Conditional Rate Lemma)**.** *Let $r \geq 1$ fix and $C > 0$ which can depend polynomially on the relevant quantities and $\varepsilon' > 0$ and let $D \geq 1$ which can depend poly-logarithmically on the relevant quantities. Further $T \subset [H]$ is a subset of time-indices with $u < t$ for all $t \in T$. Then*

$$\sum_{t\in T} \sum_{s,a\in L_k^{ut}} w_{uk}^t(s,a|s',a') \left(\frac{C(\text{llnp}(n_{tk}(s,a)) + D)}{n_{tk}(s,a)}\right)^{1/r} \leq \varepsilon' \left(\frac{\text{llnp}(n_{uk}(s',a') + D + 1}{n_{uk}(s',a')}\right)^{1/r}$$

on all but at most

$$\frac{8CAS|T|^r}{\varepsilon'^r} \text{polylog}(S, A, H, \delta^{-1}, \varepsilon'^{-1}).$$

friendly episodes E.

Proof. The proof follows mainly the structure of Lemma C.3. For the sake of completeness, we still present all steps here. Define

$$\begin{aligned}\Delta_k =& \sum_{t\in T} \sum_{s,a\in L_k^{ut}} w_{uk}^t(s,a|s',a') \left(\frac{C(\text{llnp}(n_{tk}(s,a)) + D)}{n_{tk}(s,a)}\right)^{1/r} \\ =& \sum_{t\in T} \sum_{s,a\in L_k^{ut}} w_{uk}^t(s,a|s',a')^{1-1/r} \left(w_{uk}^t(s,a|s',a') \frac{C(\text{llnp}(n_{tk}(s,a)) + D)}{n_{tk}(s,a)}\right)^{1/r}.\end{aligned}$$

We first bound using Hölder's inequality

$$\Delta_k \leq \left(\sum_{t \geq u} \sum_{s,a \in L_k^{ut}} w_{uk}^t(s,a|s',a') \frac{C|T|^{r-1}(\mathrm{llnp}(n_{tk}(s,a)) + D)}{n_{tk}(s,a)} \right)^{\frac{1}{r}}$$

Using the property in Lemma C.1 of friendly episodes as well as the fact that $w_{uk}^t(s,a|s',a') \leq 1$ and $\sum_{i<k} w_{ui}^t(s,a|s',a') \geq 4 \ln \frac{S^2A^2H^2}{\delta'} \geq 4\ln(2) \geq 2$, we bound

$$n_{tk}(s,a) \geq \frac{1}{4} n_{uk}(s',a') \sum_{i<k} w_{ui}^t(s,a|s',a') \geq \frac{1}{8} n_{uk}(s',a') \sum_{i \leq k} w_{ui}^t(s,a|s',a').$$

The function $\frac{\mathrm{llnp}(x)+D}{x}$ is monotonically decreasing in $x \geq 0$ since $D \geq 1$ (see Lemma C.6). This allows us to bound

$$\begin{aligned}
\Delta_k^r \leq & \sum_{t \in T} \sum_{s,a \in L_k^{ut}} w_{uk}^t(s,a|s',a') \frac{C|T|^{r-1}(\mathrm{llnp}(n_{tk}(s,a)) + D)}{n_{tk}(s,a)} \\
\leq & 8C|T|^{r-1} \sum_{t \in T} \sum_{s,a \in L_k^{ut}} \frac{w_{uk}^t(s,a|s',a')(\mathrm{llnp}\left(\frac{1}{8} n_{uk}(s',a') \sum_{i \leq k} w_{ui}^t(s,a|s',a')\right) + D)}{n_{uk}(s',a') \sum_{i \leq k} w_{ui}^t(s,a|s',a')} \\
\leq & 8C|T|^{r-1} \sum_{t \in T} \sum_{s,a \in L_k^{ut}} \frac{w_{uk}^t(s,a|s',a')(\mathrm{llnp}\left(\sum_{i \leq k} w_{ui}^t(s,a|s',a')\right) + \mathrm{llnp}(n_{uk}(s',a')) + D + 1)}{n_{uk}(s',a') \sum_{i \leq k} w_{ui}^t(s,a|s',a')},
\end{aligned}$$

where for the last line we used the first and last property in Lemma C.6. For notational convenience, we will use $D' = D + 1 + \mathrm{llnp}(n_{uk}(s',a'))$. Assume now $\Delta_k > \varepsilon' \left(\frac{D'}{n_{uk}(s',a')}\right)^{1/r}$. In this case the right-hand side of the inequality above is also larger than $\varepsilon'^r \left(\frac{D'}{n_{uk}(s',a')}\right)$ and there is at least one (s,a,t) with $w_{uk}^t(s,a|s',a') > w_{\min}$ and

$$\begin{aligned}
\frac{8CSA|T|^r \left(\mathrm{llnp}\left(\sum_{i \leq k} w_{ui}^t(s,a|s',a')\right) + D'\right)}{\sum_{i \leq k} w_{ui}^t(s,a|s',a')} > & D'\varepsilon'^r \\
\Leftrightarrow \frac{\left(\mathrm{llnp}\left(\sum_{i \leq k} w_{ui}^t(s,a|s',a')\right) + D'\right)}{\sum_{i \leq k} w_{ui}^t(s,a|s',a')} > & \frac{D'\varepsilon'^r}{8CSA|T|^r}.
\end{aligned}$$

Let us denote $C' = \frac{8CAS|T|^r}{\varepsilon'^r}$. Since $\frac{\mathrm{llnp}(x)+D'}{x}$ is monotonically decreasing and $x = C'^2 + 3C'$ satisfies $\frac{\mathrm{llnp}(x)+D'}{x} \leq \frac{\sqrt{x}+D'}{x} \leq D' \frac{\sqrt{x}+1}{x} \leq \frac{D'}{C'}$, we know that if $\sum_{i \leq k} w_{ui}^t(s,a|s',a') \geq C'^2 + 3C'$ then the above condition cannot be satisfied for s,a,t. Since each time the condition is satisfied, it holds that $w_{uk}^t(s,a|s',a') > w_{\min}$ and so $\sum_{i \leq k} w_{ui}^t(s,a|s',a')$ increases by at least $w_{\min}$, it can happen at most

$$m \leq \frac{AS|T|(C'^2 + 3C')}{w_{\min}}$$

times that $\Delta_k > \varepsilon' \left(\frac{D'}{n_{uk}(s',a')}\right)^{1/r}$. Define $K = \left\{k : \Delta_k > \varepsilon' \left(\frac{D'}{n_{uk}(s',a')}\right)^{1/r}\right\} \cap E$ and we know that $|K| \leq m$. Now we consider the sum

$$\begin{aligned}
\sum_{k \in K} \Delta_k^r \leq & \sum_{k \in K} 8C|T|^{r-1} \sum_{t \in T} \sum_{s,a \in L_k^{ut}} \frac{w_{uk}^t(s,a|s',a')(\mathrm{llnp}\left(\sum_{i \leq k} w_{ui}^t(s,a|s',a')\right) + D')}{n_{uk}(s',a') \sum_{i \leq k} w_{ui}^t(s,a|s',a')} \\
\leq & \frac{8C|T|^{r-1}(\mathrm{llnp}\left(C'^2 + 3C'\right) + D')}{n_{uk}(s',a')} \sum_{t \in T} \sum_{s,a \in L_k^{ut}} \sum_{k \in K} \frac{w_{uk}^t(s,a|s',a')}{\sum_{i \leq k} w_{ui}^t(s,a|s',a')} \\
\leq & \frac{8C|T|^{r-1}D'(\mathrm{llnp}\left(C'^2 + 3C'\right) + 1)}{n_{uk}(s',a')} \sum_{t \in T} \sum_{s,a \in L_k^{ut}} \sum_{k \in K} \frac{w_{uk}^t(s,a|s',a')}{\sum_{i \leq k} w_{ui}^t(s,a|s',a') \mathbb{I}\{w_{ui}^t(s,a|s',a') \geq w_{\min}\}}
\end{aligned}$$

For every (s,a,t), we consider the sequence of $w_{ui}^t(s,a|s',a') \in [w_{\min},1]$ with $i \in I = \{i \in \mathbb{N} : w_{ui}^t(s,a|s',a') \geq w_{\min}\}$ and apply Lemma C.5. This yields that

$$\sum_{k\in K} \frac{w_{uk}^t(s,a|s',a')}{\sum_{i\leq k} w_{ui}^t(s,a|s',a')\mathbb{I}\{w_{ui}^t(s,a|s',a') \geq w_{\min}\}} \leq \ln\left(\frac{me}{w_{\min}}\right)$$

and hence

$$\sum_{k\in K} \Delta_k^r \leq \frac{8CAS|T|^r D'(\operatorname{llnp}\left(C'^2+3C'\right)+1)}{n_{uk}(s',a')} \ln\left(\frac{me}{w_{\min}}\right)$$

Since each element in K has to contribute at least $\frac{D'\varepsilon'^r}{n_{uk}(s',a')}$ to this bound, we can conclude that

$$\begin{aligned}\sum_{k\in E} \mathbb{I}\{\Delta_k \geq \varepsilon'\} &= \sum_{k\in K} \mathbb{I}\{\Delta_k \geq \varepsilon'\} \\ &\leq |K| \leq \frac{8CAS|T|^r}{\varepsilon'^r} \ln\left(\frac{me}{w_{\min}}\right)\left(\operatorname{llnp}\left(C'^2+3C'\right)+1\right).\end{aligned}$$

Since $\ln\left(\frac{me}{w_{\min}}\right)\left(\operatorname{llnp}\left(C'^2+3C'\right)+1\right)$ is $\operatorname{polylog}(S,A,H,\delta^{-1},\varepsilon'^{-1})$, the proof is complete. □

Lemma C.5. *Let a_i be a sequence taking values in $[a_{\min},1]$ with $a_{\min}>0$ and $m>0$, then*

$$\sum_{k=1}^{m} \frac{a_k}{\sum_{i=1}^k a_i} \leq \ln\left(\frac{me}{a_{\min}}\right).$$

Proof. Let f be a step-function taking value a_i on $[i-1,i)$ for all i. We have $F(t) := \int_0^t f(x)dx = \sum_{i=1}^t a_i$. By the fundamental theorem of Calculus, we can bound

$$\begin{aligned}\sum_{k=1}^{m} \frac{a_k}{\sum_{i=1}^k a_i} &= \frac{a_1}{a_1} + \int_1^m \frac{f(x)}{F(x)-F(0)}dx = 1 + \ln F(m) - \ln F(1) \\ &\leq 1 + \ln(m) - \ln a_{\min} = \ln\left(\frac{me}{a_{\min}}\right),\end{aligned}$$

where the inequality follows from $a_1 \geq a_{\min}$ and $\sum_{i=1}^m a_i \leq m$. □

Lemma C.6 (Properties of llnp)**.** *The following properties hold:*

1. llnp *is continuous and nondecreasing.*
2. $f(x) = \frac{\operatorname{llnp}(nx)+D}{x}$ *with $n \geq 0$ and $D \geq 1$ is monotonically decreasing on $\mathbb{R}_+$.*
3. $\operatorname{llnp}(xy) \leq \operatorname{llnp}(x) + \operatorname{llnp}(y) + 1$ *for all $x,y \geq 0$.*

Proof. 1. For $x \leq e$ we have $\operatorname{llnp}(x) = 0$ and for $x \geq e$ we have $\operatorname{llnp}(x) = \ln(\ln(x))$ which is continuous and monotonically increasing and $\lim_{x\searrow e} \ln(\ln(x)) = 0$.

2. The function llnp is continuous as well as $1/x$ on $\mathbb{R}_+$ and therefore so it f. Further, f is differentiable except at $x = e/n$. For $x \in [0,e/n)$, we have $f(x) = D/x$ with derivative $-D/x^2 < 0$. Hence f is monotonically decreasing on $x \in [0,e/n)$. For $x > e/n$, we have $f(x) = \frac{\ln(\ln(nx))+D}{x}$ with derivative

$$-\frac{D+\ln(\ln(nx))}{x^2} + \frac{1}{x^2\ln(nx)} = \frac{1-\ln(nx)(D+\ln(\ln(nx)))}{x^2\ln(nx)}.$$

The denominator is always positive in this range so f is monotonically decreasing if and only if $\ln(nx)(D-\ln(\ln(nx))) \geq 1$. Using $D \geq 1$, we have $\ln(nx)(D+\ln(\ln(nx))) \geq 1(1+0) = 1$.

3. First note that for $xy \leq e^e$ we have $\operatorname{llnp}(xy) \leq 1 \leq \operatorname{llnp}(x) + \operatorname{llnp}(y) + 1$ and therfore the statement holds for $x, y \leq e$.

   Then consider the case that $x, y \geq e$ and $\operatorname{llnp}(x) + \operatorname{llnp}(y) + 1 - \operatorname{llnp}(xy) = \ln\ln x + \ln\ln y + 1 - \ln(\ln(x) + \ln(y)) = -\ln(a+b) + 1 + \ln(a) + \ln(b)$ where $a = \ln x \geq 1$ and $b = \ln y \geq 1$. The function $g(a, b) = -\ln(a+b) + 1 + \ln(a) + \ln(b)$ is continuous and differentiable with $\frac{\partial g}{\partial a} = \frac{b}{a(a+b)} > 0$ and $\frac{\partial g}{\partial b} = \frac{a}{b(a+b)} > 0$. Therefore, g attains its minimum on $[1, \infty) \times [1, \infty)$ at $a = 1, b = 1$. Since $g(1, 1) = 1 - \ln(2) \geq 0$, the statement also holds for $x, y \geq e$.

   Finally consider the case where $x \leq e \leq y$. Then $\operatorname{llnp}(xy) \leq \operatorname{llnp}(ey) = \ln(1 + \ln y) \leq \ln\ln y + 1 \leq \operatorname{llnp}(x) + \operatorname{llnp}(y) + 1$. Due to symmetry this also holds for $y \leq e \leq x$.

□

C.4. Decomposition of Optimality Gap

In this section we decompose the optimality gap and then bound each term individually. Finally, both rate lemmas presented in the previous section are used to determine a bound on the number of nice / friendly episodes where the optimality gap can be larger than ε. The decomposition in the following lemma is a the simpler version bounding the number of ε-suboptimal nice episodes and eventually lead to the first bound in Theorem 1.

Lemma C.7 (Optimality Gap Bound On Nice Episodes)**.** *On the good event F^c it holds that $V_1^*(s_0) - V_1^{\pi_k}(s_0) \leq \varepsilon$ on all nice episodes $k \in N$ except at most*

$$\frac{144(4 + 3H^2 + 4SH^2)ASH^2}{\varepsilon^2} \operatorname{polylog}(A, S, H, 1/\varepsilon, 1/\delta)$$

episodes.

Proof. Using optimism of the algorithm shown in Lemma C.16, we can bound

$$\begin{aligned}
&V_1^*(s_0) - V_1^{\pi_k}(s_0) \\
\leq &|\tilde{V}_1^{\pi_k}(s_0) - V_1^{\pi_k}(s_0)| \\
\leq &\sum_{t=1}^{H} \sum_{s,a} w_{tk}(s,a) |(\tilde{P}_k(s,a,t) - P(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k}| + \sum_{t=1}^{H} \sum_{s,a} w_{tk}(s,a) |\tilde{r}_k(s,a,t) - r(s,a,t)| \\
\leq &\sum_{t=1}^{H} \sum_{s,a \in L_{tk}} w_{tk}(s,a) |(\tilde{P}_k(s,a,t) - P(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k}| + \sum_{t=1}^{H} \sum_{s,a \in L_{tk}} w_{tk}(s,a) |\tilde{r}_k(s,a,t) - r(s,a,t)| \\
&+ \sum_{t=1}^{H} \sum_{s,a \notin L_{tk}} w_{tk}(s,a) |(\tilde{P}_k(s,a,t) - P(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k}| + \sum_{t=1}^{H} \sum_{s,a \notin L_{tk}} w_{tk}(s,a) |\tilde{r}_k(s,a,t) - r(s,a,t)| \\
\leq &\sum_{t=1}^{H} \sum_{s,a \notin L_{tk}} H w_{\min} + \sum_{t=1}^{H} \sum_{s,a \in L_{tk}} w_{tk}(s,a) \Big[|(\tilde{P}_k(s,a,t) - \hat{P}_k(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k}| \\
&+ |(\hat{P}_k(s,a,t) - P(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k}| + |\tilde{r}_k(s,a,t) - r(s,a,t)| \Big]
\end{aligned}$$

The first term is bounded by $c_\varepsilon \varepsilon = \frac{\varepsilon}{3}$. We now can use Lemma C.9, Lemma C.10 to bound the other terms by

$$\sum_{t=1}^{H} \sum_{s,a \in L_{tk}} w_{tk}(s,a) \sqrt{\frac{8(H + H\sqrt{S} + 2)^2}{n_{tk}(s,a)} \left(\operatorname{llnp}(n_{tk}(s,a)) + \frac{1}{2} \ln \frac{6SAH'}{\delta} \right)}.$$

We can then apply Lemma C.3 with $r = 2$, $C = 8(H + H\sqrt{S} + 2)^2$, $D = \frac{1}{2} \ln \frac{6SAH}{\delta'}$ (≥ 1 for any nontrivial setting) and $\varepsilon' = 2\varepsilon/3$ to bound this term by $\frac{2\varepsilon}{3}$ on all nice episodes except at most

$$\frac{64(H + \sqrt{S}H + 2)^2 ASH^2 3^2}{4\varepsilon^2} \operatorname{polylog}(A, S, H, 1/\varepsilon, 1/\delta) \leq \frac{144(4 + 3H^2 + 4SH^2)ASH^2}{\varepsilon^2} \operatorname{polylog}(A, S, H, 1/\varepsilon, 1/\delta)$$

Hence $V_1^*(s_0) - V_1^{\pi_k}(s_0) \leq \varepsilon$ holds on all nice episodes except those. □

The lemma below is a refined version of the bound above and uses the stronger concept of friendly episodes to eventually lead to the second bound in Theorem 1.

Lemma C.8 (Suboptimality Gap Bound On Friendly Episodes)**.** *On the good event F^c it holds that $p_0^\top(V_1^* - V_1^{\pi_k}) \leq \varepsilon$ on all friendly episodes E except at most*

$$\left(\frac{9216}{\varepsilon} + 417S\right)\frac{ASH^4}{\varepsilon}\operatorname{polylog}(S, A, H, 1/\varepsilon, \delta)$$

episodes if $\delta' \leq \frac{3AS^2H}{e^2}$.

Proof. We can further decompose the optimality gap bound in Equation C.4 in the proof of Lemma C.7 as

$$\begin{aligned}
&\sum_{t=1}^{H}\sum_{s,a\notin L_{tk}}(H+1)w_{\min} + \sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\Big[|(\tilde{P}_k(s,a,t) - \hat{P}_k(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k}| + |\tilde{r}_k(s,a,t) - r(s,a,t)| \\
&+ |(\hat{P}_k(s,a,t) - P(s,a,t))^\top V_{t+1}^*| + |(\hat{P}_k(s,a,t) - P(s,a,t))^\top (V_{t+1}^* - \tilde{V}_{t+1}^{\pi_k})|\Big]. \\
\leq& c_\varepsilon\varepsilon + \sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\Big[|(\tilde{P}_k(s,a,t) - \hat{P}_k(s,a,t))^\top \tilde{V}_{t+1}^{\pi_k}| + |\tilde{r}_k(s,a,t) - r(s,a,t)| + |(\hat{P}_k(s,a,t) - P(s,a,t))^\top V_{t+1}^*|\Big] \\
&+ \sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)|(\hat{P}_k(s,a,t) - P(s,a,t))^\top (V_{t+1}^* - \tilde{V}_{t+1}^{\pi_k})|.
\end{aligned}$$

The second term can be bounded using Lemmas C.11, C.10 and C.9 by

$$\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sqrt{\frac{32(H+1)^2}{n_{tk}(s,a)}\left(\operatorname{llnp}(n_{tk}(s,a)) + \frac{1}{2}\ln\frac{6SAH}{\delta'}\right)}.$$

which we bound by $\varepsilon/3$ using Lemma C.3 with $r = 2$, $C = 32(H+1)^2$, $D = \frac{1}{2}\ln\frac{6SAH}{\delta'}$ and $\varepsilon' = \varepsilon/3$ on all friendly episodes except at most

$$\frac{8CASH^2}{\varepsilon'^2}\operatorname{polylog}(S, A, H, 1/\varepsilon, 1/\delta) \leq \frac{9216ASH^4}{\operatorname{polylog}}(S, A, H, 1/\varepsilon, 1/\delta).$$

Finally, we apply Lemma C.12 bound to bound the last term in Equation C.4 by $\varepsilon/3$ on all friendly epsiodes but at most

$$\frac{417AS^2H^4}{\varepsilon}\operatorname{polylog}(S, A, H, 1/\delta, 1/\varepsilon).$$

It hence follows that $p_0^\top(V_1^* - V_1^{\pi_k}) \leq \varepsilon$ on all friendly episodes but at most

$$\left(\frac{9216ASH^4}{\varepsilon^2} + \frac{417AS^2H^4}{\varepsilon}\right)\operatorname{polylog}(S, A, H, 1/\delta, 1/\varepsilon).$$

□

Lemma C.9 (Algorithm Learns Fast Enough)**.** *It holds for all $s \in \mathcal{S}, a \in \mathcal{A}$ and $t \in [H]$*

$$|(\hat{P}_k(s,a,t) - \tilde{P}_k(s,a,t))^\top \tilde{V}_{t+1}| \leq \sqrt{\frac{2H^2}{n_{tk}(s,a)}\left(\operatorname{llnp}(n_{tk}(s,a)) + \frac{1}{2}\ln\frac{3SAH}{\delta'}\right)}.$$

Proof. Using the definition of the constraint in the planning step of the algorithm shown in Lemma B.1 we can bound

$$\begin{aligned}|(\hat{P}_k(s,a,t) - \tilde{P}_k(s,a,t))^\top \tilde{V}_{t+1}| \leq& \sqrt{\frac{H^2}{n_{tk}(s,a)}\left(2\operatorname{llnp}(n_{tk}(s,a)) + \ln\frac{3SAH}{\delta'}\right)}.\\ \leq& \sqrt{\frac{2H^2}{n_{tk}(s,a)}\left(\operatorname{llnp}(n_{tk}(s,a)) + \frac{1}{2}\ln\frac{3SAH}{\delta'}\right)}.\end{aligned}$$

□

Lemma C.10 (Basic Decompsition Bound)**.** *On the good event F^c it holds for all $s \in \mathcal{S}, a \in \mathcal{A}$ and $t \in [H]$*

$$\begin{aligned}|(\hat{P}_k(s,a,t) - P(s,a,t))^\top \tilde{V}_{t+1}| \leq& \sqrt{\frac{8H^2S}{n_{tk}(s,a)}\left(\operatorname{llnp}(n_{tk}(s,a)) + \frac{1}{2}\ln\frac{6SAH}{\delta'}\right)}\\ |\tilde{r}_k(s,a,t) - r(s,a,t)| \leq& \sqrt{\frac{4}{n_{tk}(s,a)}\left(\operatorname{llnp}(n_{tk}(s,a)) + \frac{1}{2}\ln\frac{3SAH}{\delta'}\right)}.\end{aligned}$$

Proof. On the good event $(F_k^{L1})c$ we have using Hölder's inequality

$$\begin{aligned}|(\hat{P}_k(s,a,t) - P(s,a,t))^\top \tilde{V}_{t+1}| \leq& \|\hat{P}_k(s,a,t) - P(s,a,t))\|_1 \|\tilde{V}_{t+1}\|_\infty\\ \leq& H\sqrt{\frac{4}{n_{tk}(s,a)}\left(2\operatorname{llnp}(n_{tk}(s,a)) + \ln\frac{3SAH(2^S-2)}{\delta'}\right)}\\ \leq& \sqrt{\frac{8H^2S}{n_{tk}(s,a)}\left(\operatorname{llnp}(n_{tk}(s,a)) + \frac{1}{2}\ln\frac{6SAH}{\delta'}\right)}.\end{aligned}$$

Further, on $(F_k^R)^c$ we have

$$\begin{aligned}|\tilde{r}_k(s,a,t) - r(s,a,t)| \leq& |\tilde{r}_k(s,a,t) - r(s,a,t)| + |\tilde{r}_k(s,a,t) - \hat{r}(s,a,t)|\\ \leq& 2\sqrt{\frac{1}{n_{tk}(s,a)}\left(2\operatorname{llnp}(n_{tk}(s,a)) + \ln\frac{3SAH}{\delta'}\right)}\end{aligned}$$

□

Lemma C.11 (Fixed V Term Confidence Bound)**.** *On the good event F^c it holds for all $s \in \mathcal{S}, a \in \mathcal{A}$ and $t \in [H]$*

$$|(\hat{P}_k(s,a,t) - P(s,a,t))^\top V^*_{t+1}| \leq \sqrt{\frac{2H^2}{n_{tk}(s,a)}\left(\operatorname{llnp} n_{tk}(s,a) + \frac{1}{2}\ln\frac{3SAH}{\delta'}\right)}$$

Proof. Since we consider the event $(F_k^V)^c$, we can bound

$$|(\hat{P}_k(s,a,t) - P(s,a,t))^\top V^*_{t+1}| \leq \sqrt{\frac{2H^2}{n_{tk}(s,a)}\left(\operatorname{llnp} n_{tk}(s,a) + \frac{1}{2}\ln\frac{3SAH}{\delta'}\right)}$$

□

Lemma C.12 (Lower Order Term)**.** *Assume $\delta' \leq \frac{3AS^2H}{e^2}$. On the good event F^c on all friendly episodes $k \in E$ except at most $\frac{417AS^2H^4}{\varepsilon}\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)$. it holds that*

$$\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)|(\hat{P}_k(s,a,t) - P(s,a,t))^\top(\tilde{V}^{\pi_k}_{t+1} - V^*_{t+1})| \leq \frac{\varepsilon}{3}.$$

Proof.

$$
\begin{aligned}
&\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)|(\hat{P}_k(s,a,t)-P(s,a,t))^\top(\tilde{V}_{t+1}^{\pi_k}-V_{t+1}^*)| \\
\leq&\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sum_{s'}\sqrt{\frac{2P(s'|s,a,t)}{n_{tk}(s,a)}\left(2\,\mathrm{llnp}(n_{tk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)}|\tilde{V}_{t+1}^{\pi_k}(s')-V_{t+1}^*(s')| \\
&+\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sum_{s'}\frac{1}{n_{tk}(s,a)}\left(2\,\mathrm{llnp}(n_{tk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)|\tilde{V}_{t+1}^{\pi_k}(s')-V_{t+1}^*(s')| \\
\leq&\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sum_{s'}\sqrt{\frac{2P(s'|s,a,t)}{n_{tk}(s,a)}\left(2\,\mathrm{llnp}(n_{tk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)\left(\tilde{V}_{t+1}^{\pi_k}(s')-V_{t+1}^*(s')\right)^2} \\
&+\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} \frac{w_{tk}(s,a)HS}{n_{tk}(s,a)}\left(2\,\mathrm{llnp}(n_{tk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right) \\
\leq&\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sqrt{\frac{2S}{n_{tk}(s,a)}\left(2\,\mathrm{llnp}(n_{tk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)P(s,a,t)^\top\left(\tilde{V}_{t+1}^{\pi_k}-V_{t+1}^*\right)^2} \\
&+\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} \frac{w_{tk}(s,a)HS}{n_{tk}(s,a)}\left(2\,\mathrm{llnp}(n_{tk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)
\end{aligned}
$$

The first inequality follows since we only consider outcomes in the event $(F_k^P)^c$, the second from the fact that value function are in the range $[0,H]$ and the third is an application of the Cauchy-Schwarz inequality. Using of optimism of the algorithm (Lemma C.16), we now bound $P(s,a,t)^\top\left(\tilde{V}_{t+1}^{\pi_k}-V_{t+1}^*\right)^2 \leq P(s,a,t)^\top\left(\tilde{V}_{t+1}^{\pi_k}-V_{t+1}^{\pi_k}\right)^2$ which we bound by $c_\varepsilon\varepsilon+\left(c_\varepsilon\varepsilon+\sqrt{\frac{C'^2}{n_{tk}(s,a)S}\left(\mathrm{llnp}(n_{tk}(s,a)+\frac{1}{2}\ln\frac{3AS^2H\varepsilon^4}{\delta'}\right)}\right)^2 \leq c_\varepsilon\varepsilon+(c_\varepsilon\varepsilon+\frac{C'}{\sqrt{S}}\sqrt{J(s,a,t)})^2$ using Lemma C.13. To keep the notation concise, we use here the shorthand $J(s,a,t)=\frac{1}{n_{tk}(s,a)}\left(\mathrm{llnp}(n_{tk}(s,a))+\frac{1}{2}\ln\frac{3e^4S^2AH}{\delta'}\right)$. This bound holds on all friendly episodes except at most $\left(32ASH^2+48AS^2H^3+AS^2H^4+16AS^2\right)\mathrm{polylog}(S,A,H,1/\delta,1/\varepsilon)$. Plugging this into the bound from above, we get the upper bound

$$
\begin{aligned}
&\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sqrt{4SJ(s,a,t)\left(c_\varepsilon\varepsilon+(c_\varepsilon\varepsilon+C'\sqrt{J(s,a,t)/S})^2\right)}+\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} 2w_{tk}(s,a)HSJ(s,a,t) \\
\leq&\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sqrt{4SJ(s,a,t)c_\varepsilon\varepsilon}+\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sqrt{4SJ(s,a,t)(c_\varepsilon\varepsilon+C'\sqrt{J(s,a,t)/S})^2} \\
&+\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} 2w_{tk}(s,a)HSJ(s,a,t) \\
=&\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)\sqrt{4(c_\varepsilon\varepsilon+c_\varepsilon^2\varepsilon^2)SJ(s,a,t)}+\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} 2w_{tk}(s,a)J(s,a,t)(C'+SH),
\end{aligned}
$$

where we used $\sqrt{a+b}\leq\sqrt{a}+\sqrt{b}$. We now bound the first term using Lemma C.3 with $r=2,\varepsilon'=\varepsilon/6,D=\frac{1}{2}\ln\frac{3e^4S^2AH}{\delta'},C=4(c_\varepsilon\varepsilon+c_\varepsilon^2\varepsilon^2)S$ on all but $\frac{8CASH^2}{\varepsilon'^2}\,\mathrm{polylog}(\dots)=\frac{192c_\varepsilon(1+c_\varepsilon\varepsilon)AS^2H^2}{\varepsilon}\,\mathrm{polylog}(\dots)$ friendly episodes by $\varepsilon/6$.

Applying Lemma C.3 with $r=1,\varepsilon'=\varepsilon/6,D=\frac{1}{2}\ln\frac{3e^4S^2AH}{\delta'}$ and $C=2(C'+HS)$, we can bound the second term by $\varepsilon/6$ on all but $\frac{8CASH}{\varepsilon'}\,\mathrm{polylog}(\dots)=\frac{96AS(C'+HS)H^2}{\varepsilon}\,\mathrm{polylog}(\dots)$ friendly episodes. Hence, it holds

$$
\sum_{t=1}^{H}\sum_{s,a\in L_{tk}} w_{tk}(s,a)|(\hat{P}_k(s,a,t)-P(s,a,t))^\top(\tilde{V}_{t+1}^{\pi_k}-V_{t+1}^*)|\leq\frac{\varepsilon}{3}
$$

on all friendly episodes except at most

$$\left(\frac{96AS(C'+HS)H^2}{\varepsilon}+\frac{192c_\varepsilon(1+c_\varepsilon\varepsilon)AS^2H^2}{\varepsilon}+32ASH^2+48AS^2H^3+AS^2H^4+16AS^2\right)\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)$$

episodes. Since $C'=\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)$, this simplifies to

$$\left(\frac{96AS}{\varepsilon}+\frac{96AS^2H^3}{\varepsilon}+\frac{64AS^2H^2}{\varepsilon}+64AS^2H^2+32ASH^2+48AS^2H^3+AS^2H^4+16AS^2\right)\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)$$

$$\leq((64+32+48+1+16)AS^2H^4+\frac{96+96+64}{\varepsilon}AS^2H^3)\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)$$

failure episodes in E. We can finally bound the failure episodes by

$$\frac{417AS^2H^4}{\varepsilon}\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon).$$

□

Lemma C.13. *On the good event F^c for any $s\in\mathcal{S}$, $a\in\mathcal{A}$ and $t\in[H]$ with $\delta'\leq\frac{3AS^2H}{e^2}$ it holds*

$$P(s,a,t)^\top(\tilde{V}_{t+1}^{\pi_k}-V_{t+1}^{\pi_k})^2\leq c_\varepsilon\varepsilon+\left(c_\varepsilon\varepsilon+\sqrt{\frac{1}{n_{tk}(s,a)S}\left(\operatorname{llnp}(n_{tk}(s,a)+\frac{1}{2}\ln\frac{3AS^2H\varepsilon^4}{\delta'}\right)}\right)^2$$

where $C'=1+\sqrt{\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}}$ on all friendly episodes except for at most

$$\left(32ASH^2+48AS^2H^3+AS^2H^4+16AS^2\right)\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)$$

episodes.

Proof. Define $L'=\{s':w_{tk}^{t+1}(s',a'|s,a)>w'_{\min}\}$ and $J(s')=\frac{\operatorname{llnp}n_{t+1k}(s',a')+\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}}{n_{t+1k}(s',a')}$ where $a'=\pi_k(s',t+1)$ and $C'=1+\sqrt{\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}}$. Using Lemma C.14, we bound

$$\begin{aligned}
&P(s,a,t)^\top(\tilde{V}_{t+1}^{\pi_k}-V_{t+1}^{\pi_k})^2=\sum_{s'}P(s'|s,a,t)(\tilde{V}_{t+1}^{\pi_k}(s')-V_{t+1}^{\pi_k}(s'))^2\\
\leq&Sw_{\min}H^2+\sum_{s'\in L'}P(s'|s,a,t)\left(c_\varepsilon\varepsilon+C'\sqrt{J(s')}\right)^2\\
\leq&c_\varepsilon\varepsilon+C'^2\sum_{s'\in L'}P(s'|s,a,t)J(s')+c_\varepsilon^2\varepsilon^2+2c_\varepsilon\varepsilon C'\sum_{s'\in L'}P(s'|s,a,t)\sqrt{J(s')}
\end{aligned}$$

on all friendly episodes except at most $\left(32+48SH+SH^2\right)ASH^2\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)$. Define now $L''=\{(s',a'):s'\in L',a'=\pi_k(s',t+1)\}$. We apply Lemma C.4 with $|T|=\{t+1\},C=1,D=\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}\geq1,r=1$ and $\varepsilon'=1/S$ to

$$\begin{aligned}
\sum_{s'\in L'}P(s'|s,a,t)J(s')=&\sum_{s',a'\in L''}\frac{w_{tk}^{t+1}(s',a'|s,a)}{n_{t+1k}(s',a')}\left(\operatorname{llnp}n_{t+1k}(s',a')+\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}\right)\\
\leq&\frac{1}{n_{tk}(s,a)S}\left(\operatorname{llnp}(n_{tk}(s,a))+\frac{1}{2}\ln\frac{3S^2AHe^4}{\delta'}\right)
\end{aligned}$$

on all but at most $8AS^2\,\mathrm{polylog}(A,S,H,1/\delta,1/\varepsilon)$ friendly episodes. Similarly, we bound

$$\begin{aligned}
&\sum_{s'\in L'} P(s'|s,a,t)\sqrt{J(s')}\\
=&\sum_{s',a'\in L''} w_{tk}^{t+1}(s',a'|s,a)\sqrt{\frac{1}{n_{t+1k}(s',a')}\left(\mathrm{llnp}\, n_{t+1k}(s',a')+\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}\right)}\\
\leq&\sqrt{\frac{1}{n_{tk}(s,a)S}\left(\mathrm{llnp}(n_{tk}(s,a))+\frac{1}{2}\ln\frac{3S^2AHe^4}{\delta'}\right)}
\end{aligned}$$

on all but at most $8AS^2\,\mathrm{polylog}(A,S,H,1/\delta,1/\varepsilon)$ friendly episodes. Hence on all friendly episodes except those failure episodes, we get

$$P(s,a,t)^\top(\tilde{V}_{t+1}^{\pi_k}-V_{t+1}^{\pi_k})^2\leq c_\varepsilon\varepsilon+\left(c_\varepsilon\varepsilon+\sqrt{\frac{C'^2}{n_{tk}(s,a)S}\left(\mathrm{llnp}(n_{tk}(s,a)+\frac{1}{2}\ln\frac{3AS^2H\varepsilon^4}{\delta'}\right)}\right)^2.$$

□

Lemma C.14. *Consider a fix $s'\in\mathcal{S}$ and $t\in[H]$, $\delta'\leq\frac{3AS^2H}{e^2}$ and the good event F^c. On all but at most*

$$\left(32+48SH+SH^2\right)ASH^2\,\mathrm{polylog}(S,A,H,1/\delta,1/\varepsilon)$$

friendly episodes E it holds that

$$V_t^{\pi_k}(s')-\tilde{V}_t^{\pi_k}(s')\leq c_\varepsilon\varepsilon+\left(1+\sqrt{12\ln\frac{3e^2S^2AH}{\delta'}}\right)\sqrt{\frac{1}{n_{tk}(s',a')}\left(\mathrm{llnp}\, n_{tk}(s',a')+\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}\right)},$$

where $a'=\pi_k(s',t)$.

Proof. For any t,s' and $a'=\pi_k(s',t)$ we use Lemma C.15 to write the value difference as

$$\tilde{V}_t^{\pi_k}(s')-V_t^{\pi_k}(s')=\sum_{u=t}^{H}\sum_{s,a}w_{tk}^u(s,a|s',a')(\tilde{P}_k(s,a,u)-P(s,a,u))^\top\tilde{V}_{u+1}+\sum_{u=t}^{H}\sum_{s,a}w_{tk}^u(s,a|s',a')(\tilde{r}_k(s,a,u)-r(s,a,u))$$

Let $L_k^{ut}=\{s,a\in\mathcal{S}\times\mathcal{A}:\,w_{tk}^u(s,a|s',a')\geq w_{\min}\}$ be the set of state-action pairs for which the conditional probability of observing is sufficiently large. Then we can bound the low-probability differences as

$$\begin{aligned}
&\sum_{u=t}^{H}\sum_{s,a\in(L_k^{ut})^c}w_{tk}^u(s,a|s',a')[(\tilde{r}_k(s,a,u)-r(s,a,u))+(P(s,a,u)-\tilde{P}_k(s,a,u))^\top\tilde{V}_{u+1}]\\
\leq&\sum_{u=t}^{H}\sum_{s,a\in(L_k^{ut})^c}w_{\min}H\leq w_{\min}H^2S=c_\varepsilon\varepsilon.
\end{aligned}$$

For the other terms with significant conditional probability, we can leverage the fact that we only consider events in $(F_k^R)^c$ and $(F_k^P)^c$ to bound

$$\begin{aligned}
&\sum_{u=t}^{H}\sum_{s,a\in L_k^{ut}}w_{tk}^u(s,a|s',a')(\tilde{r}_k(s,a,u)-r(s,a,u))\\
\leq&\sum_{u=t}^{H}\sum_{s,a\in L_k^{ut}}w_{tk}^u(s,a|s',a')\sqrt{\frac{32}{n_{tk}(s,a)}\left(\mathrm{llnp}(n_{tk}(s,a))+\frac{1}{2}\ln\frac{3SAH}{\delta'}\right)}
\end{aligned}$$

and

$$\begin{aligned}
&\sum_{u=t}^{H}\sum_{s,a\in L_k^{ut}} w_{tk}^u(s,a|s',a')(P(s,a,u)-\tilde{P}_k(s,a,u))^\top \tilde{V}_{u+1}\\
\leq&\sum_{u=t}^{H}\sum_{s,a\in L_k^{ut}} w_{tk}^u(s,a|s',a')\sum_{s''}\tilde{V}_{u+1}(s'')\sqrt{\frac{2P(s''|s,a,u)}{n_{uk}(s,a)}\left(2\operatorname{llnp}(n_{uk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)}\\
&+\sum_{u=t}^{H}\sum_{s,a\in L_k^{ut}} w_{tk}^u(s,a|s',a')\sum_{s''}\frac{\tilde{V}_{u+1}(s'')}{n_{uk}(s,a)}\left(2\operatorname{llnp}(n_{uk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)\\
\leq&\sum_{u=t}^{H}\sum_{s,a\in L_k^{ut}} w_{tk}^u(s,a|s',a')\sqrt{\frac{2SH^2}{n_{uk}(s,a)}\left(2\operatorname{llnp}(n_{uk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)}\\
&+\sum_{u=t}^{H}\sum_{s,a\in L_k^{ut}} w_{tk}^u(s,a|s',a')\frac{SH}{n_{uk}(s,a)}\left(2\operatorname{llnp}(n_{uk}(s,a))+\ln\frac{3S^2AH}{\delta'}\right)
\end{aligned}$$

where we use Cauchy Schwarz for the last inequality. Combining these individual bounds, we can upper-bound the value difference as

$$\begin{aligned}
&\tilde{V}_t^{\pi_k}(s')-V_t^{\pi_k}(s')\\
\leq&c_\varepsilon\varepsilon+\sum_{u=t}^{H}\sum_{s,a\in L_k^{ut}} w_{tk}^u(s,a|s',a')\sqrt{\frac{(4\sqrt{2}+2\sqrt{S}H)^2}{n_{uk}(s,a)}\left(\operatorname{llnp}(n_{uk}(s,a))+\frac{1}{2}\ln\frac{3S^2AH}{\delta'}\right)}\\
&+\sum_{u=t}^{H}\sum_{s,a\in L_k^{ut}} w_{tk}^u(s,a|s',a')\frac{2SH}{n_{uk}(s,a)}\left(\operatorname{llnp}(n_{uk}(s,a))+\frac{1}{2}\ln\frac{3S^2AH}{\delta'}\right) \qquad (12)
\end{aligned}$$

We now apply Lemma C.4 with $r=2, D=\frac{1}{2}\ln\frac{3S^2AH}{\delta'}, C=(4\sqrt{2}+2\sqrt{S}H)^2, T=\{t+1,t+2,\dots H\}$ and $\varepsilon'=1$ and get that the second term above is bounded by

$$\sqrt{\frac{1}{n_{tk}(s',a')}\left(\operatorname{llnp} n_{tk}(s',a')+\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}\right)}$$

on all friendly episodes but at most

$$\frac{8CASH^2}{\varepsilon'^2}\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)=(32+16\sqrt{2S}H+SH^2)ASH^2\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)$$

episodes. We apply Lemma C.4 again to the final term in Equation (12) above with $r=1, D=\frac{1}{2}\ln\frac{3S^2AH}{\delta'}\geq 1, T=\{t+1,t+2,\dots H\}, C=2SH$ and $\varepsilon'=1$. Then the final term is bounded by $\frac{1}{n_{tk}(s',a')}\left(\operatorname{llnp} n_{tk}(s',a')+\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}\right)$. on all friendly episodes but

$$\frac{8CASH^2}{\varepsilon'^2}\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)=16AS^2H^3\operatorname{polylog}(S,A,H,1/\delta,1/\varepsilon)$$

many. Combining these bounds, we arrive at

$$\begin{aligned}
&V_t^{\pi_k}(s')-\tilde{V}_t^{\pi_k}(s')\\
\leq&c_\varepsilon\varepsilon+\sqrt{\frac{1}{n_{tk}(s',a')}\left(\operatorname{llnp} n_{tk}(s',a')+\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}\right)}\\
&+\frac{1}{n_{tk}(s',a')}\left(\operatorname{llnp} n_{tk}(s',a')+\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}\right)
\end{aligned}$$

$$\leq c_\varepsilon \varepsilon + \left(1 + \sqrt{\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}}\right)\sqrt{\frac{1}{n_{tk}(s',a')}\left(\operatorname{llnp} n_{tk}(s',a') + \frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}\right)},$$

where we bounded $\sqrt{\frac{1}{n_{tk}(s',a')}\left(\operatorname{llnp} n_{tk}(s',a') + \frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}\right)}$ by $\frac{1}{2}\ln\frac{3e^2S^2AH}{\delta'}$ since it is decreasing in $n_{tk}(s',a')$ and we therefore can simply use $n_{tk}(s',a') = 1$ (entire bound holds trivially for $n_{tk}(s',a') = 0$). □

C.5. Useful Lemmas

Lemma C.15 (Value Difference Lemma)**.** *For any two MDPs M' and M'' with rewards r' and r'' and transition probabilities P' and P'', the difference in values with respect to the same policy π can be written as*

$$V_i'(s) - V_i''(s) = \mathbb{E}''\left[\sum_{t=i}^{H}(r'(s_t,a_t,t) - r''(s_t,a_t,t))\middle| s_i = s\right] + \mathbb{E}''\left[\sum_{t=i}^{H}(P'(s_t,a_t,t) - P''(s_t,a_t,t))^\top V_{t+1}'\middle| s_i = s\right]$$

where $V_{H+1}' = V_{H+1}'' = \vec{0}$ and the expectation $\mathbb{E}'$ is taken w.r.t to P' and π and $\mathbb{E}''$ w.r.t. P'' and π.

Proof. For $i = H+1$ the statement is trivially true. We assume now it holds for $i+1$ and show it holds also for i. Using only this induction hypothesis and basic algebra, we can write

$$\begin{aligned}
&V_i'(s) - V_i''(s)\\
=&\mathbb{E}_\pi[r'(s_i,a_i,i) + {V_{i+1}'}^\top P'(s_i,a_i,i) - r''(s_i,a_i,i) - {V_{i+1}''}^\top P''(s_i,a_i,i)|s_i = s]\\
=&\mathbb{E}_\pi[r'(s_i,a_i,i) - r''(s_i,a_i,i)|s_i = s] + \mathbb{E}_\pi\left[\sum_{s'\in\mathcal{S}} V_{i+1}'(s')(P'(s'|s_i,a_i,i) - P''(s'|s_i,a_i,i))\middle| s_i = s\right]\\
&+ \mathbb{E}_\pi\left[\sum_{s'\in\mathcal{S}} P''(s'|s_i,a_i,i)(V_{i+1}'(s') - V_{i+1}''(s'))\middle| s_i = s\right]\\
=&\mathbb{E}_\pi[r'(s_i,a_i,i) - r''(s_i,a_i,i)|s_i = s] + \mathbb{E}_\pi\left[\sum_{s'\in\mathcal{S}} V_{i+1}'(s')(P'(s'|s_i,a_i,i) - P''(s'|s_i,a_i,i))\middle| s_i = s\right]\\
&+ \mathbb{E}''\left[V_{i+1}'(s_{i+1}) - V_{i+1}''(s_{i+1}))\middle| s_i = s\right]\\
=&\mathbb{E}_\pi[r'(s_i,a_i,i) - r''(s_i,a_i,i)|s_i = s] + \mathbb{E}_\pi\left[\sum_{s'\in\mathcal{S}} V_{i+1}'(s')(P'(s'|s_i,a_i,i) - P''(s'|s_i,a_i,i))\middle| s_i = s\right]\\
&+ \mathbb{E}''\left[\mathbb{E}''\left[\sum_{t=i+1}^{H}(r'(s_t,a_t,t) - r''(s_t,a_t,t))\middle| s_{i+1}\right] + \mathbb{E}''\left[\sum_{t=i+1}^{H}(P'(s_t,a_t,t) - P''(s_t,a_t,t))^\top V_{t+1}'\middle| s_{i+1}\right]\middle| s_i = s\right]\\
=&\mathbb{E}''\left[\sum_{t=i}^{H}(r'(s_t,a_t,t) - r''(s_t,a_t,t))\middle| s_i = s\right] + \mathbb{E}''\left[\sum_{t=i}^{H}(P'(s_t,a_t,t) - P''(s_t,a_t,t))^\top V_{t+1}'\middle| s_i = s\right]
\end{aligned}$$

where the last equality follows from law of total expectation □

Lemma C.16 (Algorithm ensures optimism)**.** *On the good event F^c it holds that for all episodes k, $t \in [H]$, $s \in \mathcal{S}$ that*

$$V_t^{\pi_k}(s) \leq V_t^*(s) \leq \tilde{V}_t^{\pi_k}(s).$$

Proof. The first inequality follows simply from the definition of the optimal value function V^*.

Since all outcome we consider are in the event $(F_k^V)^c$, we know that the true transition probabilities P, the optimal policy π^* and optimal policy V^* are a feasible solution for the optimistic planning problem in Lemma B.1 that UBEV solves. It therefore follows immediately that $p_0^\top \tilde{V}_1^{\pi_k} \geq p_0^\top V_1^*$. □

D. Dependency of optimistic values and empirical transition probabilities

In this section we shed more light on why the dependency between optimistic state values $\tilde{V}_{t+1}$ and empirical transition probabilities $\hat{P}_k(s,a,t)$ poses a major challenge for obtaining linear-polylog(S) sample complexity bounds.

We illustrate the issue using the example of a simple two-timestep MDP with just one action ($H=2, A=1$) and S states. The agent starts always in state 1 and transitions with equal probability to any state, i.e., $P(\cdot|s=1,a=1,t=1)=\frac{1}{S}$. The agent always received reward of 0 everywhere. One popular way to ensure optimism is to simply add an exploration bonus to empirical rewards when computing the optimal policy that scales with rate $1/\sqrt{n}$ with the number of observations for the considered state. In this example, the optimistic value $\tilde{V}_2(s)=\frac{C}{\sqrt{n(s)}}$ for any state s at time 2 is then simply a constant C divided by the square-root of the state visits $n(s)$. Assume now sample a fixed number of N episodes from this MDP. Then the difference in expected next state value which we want to bound by concentration results are

$$(\hat{P}(s=1,a=1,t=1)-P(s=1,a=1,t=1))^\top \tilde{V}_2 = \sum_{s\in\mathcal{S}}\left(\frac{n(s)}{N}-\frac{1}{S}\right)\frac{C}{\sqrt{n(s)}},$$

where $n(s)$ is the (random) number of times state s has been observed at time $t=2$. The marginal distribution of each $n(s)$ is a Multinomial($1/S$, N). The expected value of the random quantity above is therefore

$$\mathbb{E}[(\hat{P}(s=1,a=1,t=1)-P(s=1,a=1,t=1))^\top \tilde{V}_2] = \frac{CS\mathbb{E}[\sqrt{m}]}{N} - C\mathbb{E}[m^{-1/2}]$$

where $m\sim \text{Multinomial}(1/S,N)$. Applying now Jensen's ineuqlity to the first and second term separately, we get that

$$\mathbb{E}[(\hat{P}(s=1,a=1,t=1)-P(s=1,a=1,t=1))^\top \tilde{V}_2] < \frac{CS\sqrt{N}}{N\sqrt{S}} - C\frac{\sqrt{S}}{\sqrt{N}} = 0$$

due to the strict concavity of $\sqrt{\cdot}$ and strict convexity of $(\cdot)^{-1/2}$. Hence, the random exploration bonus applied to the rewards introduces a bias and therefore standard concentration results are not directly applicable. Intuitively, this bias occurs because overestimation of the transition probability to a particular state corresponds to a higher than expected number of observations of that state at time $t=2$ and therefore its optimistic value is low. Therefore there is a correlation between estimation error and optimistic value across successor states. While this simple example uses exploration bonuses applied to rewards for optimism this correlation is a general issue of optimistic values and is also present in the form of optimism used by UBEV. We can therefore not directly apply concentration results to such terms.

Indeed, unfortunately this issue causes several prior sample complexity results to be invalid (Pazis & Parr (2016) and Pazis et al. (2016)). The key issue is that it is appealing to derive bounds of the above form on the product of a value function with the difference in transition estimates that assume the value is a fixed quantity, as has been done in that work. Unfortunately, while such bounds are valid, they are then applied in algorithms where the value function is computed using samples related (either directly or through the policy followed) to the samples used to estimate the transition models. Our above analysis demonstrates that the resulting estimator is biased. Indeed, while one might wonder if such a result is primarily of theoretical interest, indicating some looseness in the existing analysis, our empirical simulations also showed a significant non-zero correlation between the number of states and the resulting confidence bound derived assuming the value function is a fixed quantity, compared to the desired result which is that the confidence bound is independent of the state space size. This highlights the need for alternative analysis to address this issue.

E. Conversion Theorem for Uniform-PAC to Regret

Theorem E.1 (Uniform-PAC to Regret Conversion Theorem)**.** *Assume on some event E an algorithm follows for all ε an ε-optimal policy π_k, i.e., $V_1^*(s_0)-V^{\pi_k}(s_0)\leq\varepsilon$, on all but at most*

$$\frac{C_1}{\varepsilon}\left(\ln\frac{C_3}{\varepsilon}\right)^k + \frac{C_2}{\varepsilon^2}\left(\ln\frac{C_3}{\varepsilon}\right)^{2k}$$

episodes where $C_1\geq C_2\geq 2$ and $C_3\geq\max\{H,e\}$ and C_1,C_2,C_3 do not depend on ε . Then this algorithm has on this event a regret of

$$R(T)\leq(\sqrt{C_2T}+C_1)\,\text{polylog}(T,C_3,C_1)=O(\sqrt{C_2T}\,\text{polylog}(T,C_3,C_1,H))$$

for all number of episodes T.

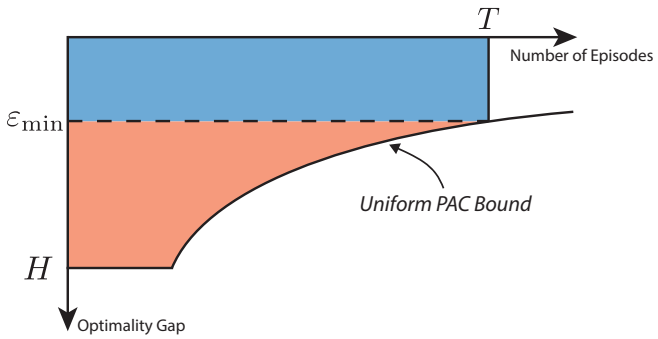

Figure 2: Relation of PAC-bound and Regret; The area of the shaded regions are a bound on the regret after T episodes.

Proof. The mistake bound $g(\varepsilon) = \frac{C_1}{\varepsilon}\left(\ln\frac{C_3}{\varepsilon}\right)^k + \frac{C_2}{\varepsilon^2}\left(\ln\frac{C_3}{\varepsilon}\right)^{2k} \leq T$ is monotonically decreasing for $\varepsilon \in (0, H]$. For a given T large enough, we can therefore find an $\varepsilon_{\min} \in (0, H]$ such that $g(\varepsilon) \leq T$ for all $\varepsilon \in (\varepsilon_{\min}, H]$. The regret $R(T)$ of the algorithm can then be bounded as follows

$$R(T) \leq T\varepsilon_{\min} + \int_{\varepsilon_{\min}}^{H} g(\varepsilon)d\varepsilon.$$

This bound assumes the worst case where first the algorithm makes the worst mistakes possible with regret H and subsequently less and less severe mistakes controlled by the mistake bound. For a better intuition, see Figure 2.

We first find a suitable $\varepsilon_{\min}$. Define $y = \frac{1}{\varepsilon}\left(\ln\frac{C_3}{\varepsilon}\right)^k$ then since g is monotonically decreasing, it is sufficient to find a ε with $g(\varepsilon) \leq T$. That is equivalent to $C_1 y + C_2 y^2 \leq T$ for which

$$\frac{1}{\varepsilon}\left(\ln\frac{C_3}{\varepsilon}\right)^k = y \leq \frac{C_1}{2C_2} + \frac{\sqrt{C_1^2 + 4TC_2}}{2C_2} =: a$$

is sufficient. We set now

$$\varepsilon_{\min} = \frac{\ln(C_3 a)^k}{a} = \frac{2C_2}{C_1 + \sqrt{C_1^2 + 4TC_2}}\left(\ln\frac{(C_1 + \sqrt{C_1^2 + 4TC_2})C_3}{2C_2}\right)^k$$

which is a valid choice as

$$\begin{aligned}\frac{1}{\varepsilon_{\min}}\left(\ln\frac{C_3}{\varepsilon_{\min}}\right)^k &= \frac{a}{\ln(C_3 a)^k}\left(\ln\frac{C_3 a}{\ln(C_3 a)^k}\right)^k = \frac{a}{\ln(C_3 a)^k}\left(\ln(C_3 a) - k\ln\ln(C_3 a)\right)^k \\ &\leq \frac{a}{\ln(C_3 a)^k}\left(\ln(C_3 a)\right)^k = a.\end{aligned}$$

We now first bound the regret further as

$$\begin{aligned}R(T) \leq& T\varepsilon_{\min} + \int_{\varepsilon_{\min}}^{H} g(\varepsilon)d\varepsilon \leq T\varepsilon_{\min} + C_1\left(\ln\frac{C_3}{\varepsilon_{\min}}\right)^k \int_{\varepsilon_{\min}}^{H}\frac{1}{\varepsilon}d\varepsilon + C_2\left(\ln\frac{C_3}{\varepsilon_{\min}}\right)^{2k}\int_{\varepsilon_{\min}}^{H}\frac{1}{\varepsilon^2}d\varepsilon \\ =& T\varepsilon_{\min} + C_1\left(\ln\frac{C_3}{\varepsilon_{\min}}\right)^k \ln\frac{H}{\varepsilon_{\min}} + C_2\left(\ln\frac{C_3}{\varepsilon_{\min}}\right)^{2k}\left[\frac{1}{\varepsilon_{\min}} - \frac{1}{H}\right]\end{aligned}$$

and then use the choice of $\varepsilon_{\min}$ from above to look at each of the terms in this bound individually. In the following bounds we extensively use the fact $\ln(a+b) \leq \ln(a) + \ln(b) = \ln(ab)$ for all $a, b \geq 2$ and that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ which holds

for all $a, b \geq 0$.

$$\begin{aligned} T\varepsilon_{\min} &= \frac{2TC_2}{C_1 + \sqrt{C_1^2 + 4TC_2}} \left(\ln \frac{C_3(C_1 + \sqrt{C_1^2 + 4TC_2})}{2C_2} \right)^k \\ &\leq \frac{2TC_2}{\sqrt{4TC_2}} \left(\ln C_3 + \ln C_1 + \ln C_1 + \ln \frac{2\sqrt{TC_2}}{2C_2} \right)^k \\ &\leq \sqrt{TC_2} \left(\ln(C_3 C_1^2 \sqrt{T}) \right)^k \end{aligned}$$

Now for a $C \geq 0$ we first look at

$$\begin{aligned} \ln \frac{C}{\varepsilon_{\min}} &= \ln C + \ln \frac{C_1 + \sqrt{C_1^2 + 4TC_2}}{2C_2} - k \ln \ln \frac{C_3(C_1 + \sqrt{C_1^2 + 4TC_2})}{2C_2} \\ &\leq \ln C + \ln \frac{C_1 + \sqrt{C_1^2 + 4TC_2}}{2C_2} \\ &\leq \ln C + \ln C_1 + \ln C_1 + \ln \frac{\sqrt{4TC_2}}{2C_2} \\ &\leq \ln(C C_1^2 \sqrt{T}) \end{aligned}$$

where the first inequality follows from the fact that $\frac{C_3(C_1+\sqrt{C_1^2+4TC_2})}{2C_2} \geq \frac{C_3 2C_1}{2C_2} \geq e$. Hence, we can bound

$$C_1 \left(\ln \frac{C_3}{\varepsilon_{\min}} \right)^k \ln \frac{H}{\varepsilon_{\min}} \leq C_1 \left(\ln(C_3 C_1^2 \sqrt{T}) \right)^k \ln(H C_1^2 \sqrt{T}).$$

Now since

$$\frac{1}{\varepsilon_{\min}} = \frac{C_1 + \sqrt{C_1^2 + 4TC_2}}{2C_2} \left(\ln \frac{C_3(C_1 + \sqrt{C_1^2 + 4TC_2})}{2C_2} \right)^{-k} \leq \frac{C_1}{C_2} + \sqrt{\frac{T}{C_2}}$$

we get

$$\begin{aligned} C_2 \left(\ln \frac{C_3}{\varepsilon_{\min}} \right)^{2k} \left[\frac{1}{\varepsilon_{\min}} - \frac{1}{H} \right] &\leq C_2 \left(\ln(C_3 C_1^2 \sqrt{T}) \right)^{2k} \left[\frac{C_1}{C_2} + \sqrt{\frac{T}{C_2}} \right] \\ &\leq \left(\ln(C_3 C_1^2 \sqrt{T}) \right)^{2k} \left[C_1 + \sqrt{TC_2} \right]. \end{aligned}$$

As a result we can conclude that $R(T) \leq (\sqrt{C_2 T} + C_1) \operatorname{polylog}(T, C_3, C_1, H) = O(\sqrt{C_2 T} \operatorname{polylog}(T, C_3, C_1, H))$. □

F. General Concentration Bounds

Lemma F.1. *Let $X_1, X_2, \ldots$ be a martingale difference sequence adapted to filtration $\{\mathcal{F}_t\}_{t=1}^{\infty}$ with X_t conditionally σ^2-subgaussian so that $\mathbb{E}[\exp(\lambda(X_t - \mu))|\mathcal{F}_{t-1}] \leq \exp(\lambda^2\sigma^2/2)$ almost surely for all $\lambda \in \mathbb{R}$. Then with $\hat{\mu}_t = \frac{1}{t}\sum_{i=1}^t X_i$ we have for all $\delta \in (0, 1]$*

$$\mathbb{P}\left(\exists t : |\hat{\mu}_t - \mu| \geq \sqrt{\frac{4\sigma^2}{t} \left(2 \operatorname{llnp}(t) + \ln \frac{3}{\delta} \right)} \right) \leq 2\delta \,.$$

Proof. Let $S_t = \sum_{s=1}^{t}(X_s - \mu)$. Then

$$\begin{aligned}
&\mathbb{P}\left(\exists t : \hat{\mu}_t - \mu \geq \sqrt{\frac{4\sigma^2}{t}\left(2\,\mathrm{llnp}(t) + \ln\frac{3}{\delta}\right)}\right)\\
\leq&\mathbb{P}\left(\exists t : S_t \geq \sqrt{4\sigma^2 t\left(2\,\mathrm{llnp}(t) + \ln\frac{3}{\delta}\right)}\right)\\
\leq&\sum_{k=0}^{\infty}\mathbb{P}\left(\exists t \in [2^k, 2^{k+1}] : S_t \geq \sqrt{4\sigma^2 t\left(2\,\mathrm{llnp}(t) + \ln\frac{3}{\delta}\right)}\right)\\
\leq&\sum_{k=0}^{\infty}\mathbb{P}\left(\exists t \leq 2^{k+1} : S_t \geq \sqrt{2\sigma^2 2^{k+1}\left(2\,\mathrm{llnp}(2^k) + \ln\frac{3}{\delta}\right)}\right)
\end{aligned}$$

We now consider $M_t = \exp(\lambda S_t)$ for $\lambda > 0$ which is a nonnegative sub-martingale and use the short-hand $f = \sqrt{2\sigma^2 2^{k+1}\left(2\,\mathrm{llnp}(2^k) + \ln\frac{3}{\delta}\right)}$. Then by Doob's maximal inequality for nonnegative submartingales

$$\mathbb{P}\left(\exists t \leq 2^{k+1} : S_t \geq f\right) = \mathbb{P}\left(\max_{t\leq 2^{k+1}} M_t \geq \exp(\lambda f)\right) \leq \frac{\mathbb{E}[M_{2^{k+1}}]}{\exp(\lambda f)} \leq \exp\left(2^{k+1}\frac{\lambda^2\sigma^2}{2} - \lambda f\right).$$

Choosing the optimal $\lambda = \frac{f}{\sigma^2 2^{k+1}}$ we obtain the bound

$$\begin{aligned}
\mathbb{P}\left(\exists t \leq 2^{k+1} : S_t \geq f\right) \leq&\exp\left(-\frac{f^2}{2^{k+2}\sigma^2}\right) = \exp\left(-2\,\mathrm{llnp}(2^k) - \ln\frac{3}{\delta}\right) = \frac{\delta}{3}\exp\left(-2\,\mathrm{llnp}(2^k)\right) \qquad (13)\\
=&\frac{\delta}{3}\exp\left(-\max\{0, 2\ln\max\{0, \ln 2^k\}\}\right) = \frac{\delta}{3}\min\left\{1, (k\ln 2)^{-2}\right\}\\
\leq&\frac{\delta}{3}\min\left\{1, \frac{1}{k^2\ln 2}\right\}.
\end{aligned}$$

Plugging this back in the bound from above, we get

$$\begin{aligned}
\mathbb{P}\left(\exists t : \hat{\mu}_t - \mu \geq \sqrt{\frac{4\sigma^2}{t}\left(2\,\mathrm{llnp}(t) + \ln\frac{3}{\delta}\right)}\right) \leq&\frac{\delta}{3}\sum_{k=0}^{\infty}\min\left\{1, \frac{1}{k^2\ln(2)}\right\}\\
=&\delta\,\frac{1}{3}\left(\frac{\pi^2}{6\ln 2} + 2 - 1/\ln(2)\right) \leq \delta\,. \qquad (14)
\end{aligned}$$

For the other side, the argument follows completely analogously with

$$\begin{aligned}
\mathbb{P}\left(\exists t \leq 2^{k+1} : S_t \leq -f\right) =&\mathbb{P}\left(\exists t \leq 2^{k+1} : -S_t \geq f\right)\\
=&\mathbb{P}\left(\max_{t\leq 2^{k+1}}\exp(-\lambda S_t) \geq \exp(\lambda f)\right)\\
\leq&\frac{\mathbb{E}[\exp(-\lambda S_{2^{k+1}})]}{\exp(\lambda f)} \leq \exp\left(2^{k+1}\frac{\lambda^2\sigma^2}{2} - \lambda f\right).
\end{aligned}$$

□

Lemma F.2. *Let $X_1, X_2, \ldots$ be a sequence of Bernoulli random variables with bias $\mu \in [0,1]$. Then for all $\delta \in (0,1]$*

$$\mathbb{P}\left(\exists t : |\hat{\mu}_t - \mu| \geq \sqrt{\frac{2\mu}{t}\left(2\,\mathrm{llnp}(t) + \ln\frac{3}{\delta}\right)} + \frac{1}{t}\left(2\,\mathrm{llnp}(t) + \ln\frac{3}{\delta}\right)\right) \leq 2\delta$$

Proof.

$$\begin{aligned}
&\mathbb{P}\left(\exists t : \hat{\mu}_t - \mu \geq \sqrt{\frac{2\mu}{t}\left(2\operatorname{llnp}(t) + \ln\frac{3}{\delta}\right)} + \frac{1}{t}\left(2\operatorname{llnp}(t) + \ln\frac{3}{\delta}\right)\right)\\
=&\mathbb{P}\left(\exists t : S_t \geq \sqrt{2\mu t\left(2\operatorname{llnp}(t) + \ln\frac{3}{\delta}\right)} + 2\operatorname{llnp}(t) + \ln\frac{3}{\delta}\right)\\
\leq&\sum_{k=0}^{\infty}\mathbb{P}\left(\exists t \leq 2^{k+1} : S_t \geq \sqrt{2\mu 2^k\left(2\operatorname{llnp}(2^k) + \ln\frac{3}{\delta}\right)} + 2\operatorname{llnp}(2^k) + \ln\frac{3}{\delta}\right)
\end{aligned}$$

Let $g = 2\operatorname{llnp}(2^k) + \ln\frac{3}{\delta}$ and $f = \sqrt{2^{k+1}\mu g} + g$. Further define $S_t = \sum_{i=1}^{t} X_i - t\mu$ and $M_t = \exp(\lambda S_t)$ which is by construction a nonnegative submartingale. Applying Doob's maximal inequality for nonnegative submartingales, we bound

$$\mathbb{P}\left(\exists t \leq 2^{k+1} : S_t \geq f\right) = \mathbb{P}\left(\max_{i \leq 2^{k+1}} M_i \geq \exp(\lambda f)\right) \leq \frac{\mathbb{E}[M_{2^{k+1}}]}{\exp(\lambda f)} = \exp\left(\ln \mathbb{E}[M_{2^{k+1}}] - \lambda f\right).$$

Since this holds for all $\lambda \in \mathbb{R}$, we can bound

$$\mathbb{P}\left(\exists t \leq 2^{k+1} : S_t \geq f\right) \leq \exp\left(-\sup_{\lambda \in \mathbb{R}}\left(\lambda f - \ln \mathbb{E}[M_{2^{k+1}}]\right)\right)$$

and using Corollary 2.11 by Boucheron et al. (2013) (see also note below proof of Corollary 2.11) bound that by

$$\exp\left(-\frac{f^2}{2(2^{k+1}\mu + f/3)}\right)$$

We now argue that this quantity can be upper-bounded by $\exp(-g)$. This is equivalent to

$$\begin{aligned}
-\frac{f^2}{2(2^{k+1}\mu + f/3)} \leq& -g\\
f^2 \geq& 2g(2^{k+1}\mu + f/3) = \frac{2}{3}gf + \frac{2^{k+2}}{3}\mu g\\
g^2 + 2\sqrt{2^{k+1}\mu g}g + 2^{k+1}\mu g \geq& \frac{2}{3}g^2 + \frac{2}{3}\sqrt{2^{k+1}\mu g}g + \frac{2^{k+2}}{3}\mu g\\
\frac{1}{3}g^2 + \frac{4}{3}\sqrt{2^{k+1}\mu g}g + \frac{1}{3}2^{k+1}\mu g \geq& 0.
\end{aligned}$$

Each line is an equivalent inequality since $g, f \geq 0$ and each term on the left in the final inequality is nonnegative. Hence, we get $\mathbb{P}\left(\exists t \leq 2^{k+1} : S_t \geq f\right) \leq \exp(-g)$. Following now the arguments from the proof of Lemma F.1 in Equations (13)–(14), we obtain that

$$\mathbb{P}\left(\exists t : \hat{\mu}_t - \mu \geq \sqrt{\frac{2\mu}{t}\left(2\operatorname{llnp}(t) + \ln\frac{3}{\delta}\right)} + \frac{1}{t}\left(2\operatorname{llnp}(t) + \ln\frac{3}{\delta}\right)\right) \leq \delta.$$

For the other direction, we proceed analogously to above and arrive at

$$\mathbb{P}\left(\exists t \leq 2^{k+1} : -S_t \geq f\right) \leq \exp\left(-\sup_{\lambda \in \mathbb{R}}\left(-\lambda f - \ln \mathbb{E}[M_{2^{k+1}}]\right)\right)$$

which we bound similarly to above by

$$\exp\left(-\frac{f^2}{2(2^{k+1}\mu - f/3)}\right) \leq \exp\left(-\frac{f^2}{2(2^{k+1}\mu + f/3)}\right) \leq \exp(-g).$$

□

Lemma F.3 (Uniform L1-Deviation Bound for Empirical Distribution)**.** *Let $X_1, X_2, \ldots$ be a sequence of i.i.d. categorical variables on $[U]$ with distribution P. Then for all $\delta \in (0,1]$*

$$\mathbb{P}\left(\exists t \,:\, \|\hat{P}_t - P\|_1 \geq \sqrt{\frac{4}{t}\left(2\,\mathrm{llnp}(t) + \ln\frac{3(2^U-2)}{\delta}\right)}\right) \leq \delta$$

where $\hat{P}_t$ is the empirical distribution based on samples $X_1 \ldots X_t$.

Proof. We use the identity $\|Q-P\|_1 = 2\max_{B\subseteq\mathcal{B}} Q(B) - P(B)$ which holds for all distributions P, Q defined on the finite set $\mathcal{B}$ to bound

$$\begin{aligned}
&\mathbb{P}\left(\exists t \,:\, \|\hat{P}_t - P\|_1 \geq \sqrt{\frac{4}{t}\left(2\,\mathrm{llnp}(t) + \ln\frac{3(2^U-2)}{\delta}\right)}\right)\\
=&\mathbb{P}\left(\max_{t,B\subseteq[U]} \hat{P}_t(B) - P(B) \geq \frac{1}{2}\sqrt{\frac{4}{t}\left(2\,\mathrm{llnp}(t) + \ln\frac{3(2^U-2)}{\delta}\right)}\right)\\
\leq& \sum_{B\subseteq[U]} \mathbb{P}\left(\max_t \hat{P}_t(B) - P(B) \geq \sqrt{\frac{1}{t}\left(2\,\mathrm{llnp}(t) + \ln\frac{3(2^U-2)}{\delta}\right)}\right).
\end{aligned}$$

Define now $S_t = \sum_{i=1}^t \mathbb{I}\{X_1 \in B\} - tP(B)$ which is a martingale sequence. Then the last line above is equivalent to

$$\begin{aligned}
&\sum_{B\subseteq[U]} \mathbb{P}\left(\max_t S_t \geq \sqrt{t\left(2\,\mathrm{llnp}(t) + \ln\frac{3(2^U-2)}{\delta}\right)}\right)\\
\leq& \sum_{B\subseteq[U]} \mathbb{P}\left(\max_{k\in\mathbb{N},t\in[2^k,2^{k+1}]} S_t \geq \sqrt{t\left(2\,\mathrm{llnp}(t) + \ln\frac{3(2^U-2)}{\delta}\right)}\right)\\
\leq& \sum_{B\subseteq[U]}\sum_{k=0}^{\infty} \mathbb{P}\left(\max_{t\in[2^k,2^{k+1}]} S_t \geq \sqrt{t\left(2\,\mathrm{llnp}(t) + \ln\frac{3(2^U-2)}{\delta}\right)}\right)\\
\leq& \sum_{B\subseteq[U]}\sum_{k=0}^{\infty} \mathbb{P}\left(\max_{t\leq 2^{k+1}} S_t \geq \sqrt{2^k\left(2\,\mathrm{llnp}(2^k) + \ln\frac{3(2^U-2)}{\delta}\right)}\right)\\
=& \sum_{B\subseteq[U]}\sum_{k=0}^{\infty} \mathbb{P}\left(\max_{t\leq 2^{k+1}} \exp(\lambda S_t) \geq \exp(\lambda f)\right)\\
=& \sum_{B\subseteq[U],B\neq\emptyset,B\neq[U]}\sum_{k=0}^{\infty} \mathbb{P}\left(\max_{t\leq 2^{k+1}} \exp(\lambda S_t) \geq \exp(\lambda f)\right)
\end{aligned}$$

where $f = \sqrt{2^k\left(2\,\mathrm{llnp}(2^k) + \ln\frac{3(2^U-2)}{\delta}\right)}$ and $\lambda \in \mathbb{R}$ and the last equality follows from the fact that for $B = \emptyset$ and $B = [U]$ the difference between the distributions has to be 0. Since $\mathbb{I}\{X_1 \in B\} - tP(B)$ is a centered Bernoulli variable it is $1/2$-subgaussian and so S_t satisfies $\mathbb{E}[\exp(\lambda S_t)] \leq \exp(\lambda^2 t/8)]$. Since S_t is a martingale, $\exp(\lambda S_t)$ is a nonnegative sub-martingale and we can apply the maximal inequality to bound

$$\mathbb{P}\left(\max_{t\leq 2^{k+1}} \exp(\lambda S_t) \geq \exp(\lambda f)\right) \leq \exp\left(\frac{1}{8}\lambda^2 2^{k+1} - \lambda f\right).$$

Choosing $\lambda = \frac{4f}{2^{k+1}}$, we get $\mathbb{P}\left(\max_{t\leq 2^{k+1}} \exp(\lambda S_t) \geq \exp(\lambda f)\right) \leq \exp\left(-\frac{f^2}{2^k}\right)$. Hence, using the same steps as in the

proof of Lemma F.1, we get $\mathbb{P}\left(\max_{t\leq 2^{k+1}} \exp(\lambda S_t) \geq \exp(\lambda f)\right) \leq \frac{\delta}{3(2^{[U]}-2)} \min\left\{1, \frac{1}{k^2 \ln 2}\right\}$ and then

$$\mathbb{P}\left(\exists t : \|\hat{P}_t - P\|_1 \geq \sqrt{\frac{4}{t}\left(2\,\mathrm{llnp}(t) + \ln\frac{3(2^U-2)}{\delta}\right)}\right)$$
$$\leq \sum_{B\subseteq[U], B\neq\emptyset, B\neq[U]} \frac{\delta}{3(2^{[U]}-2)} \sum_{k=0}^{\infty} \min\left\{1, \frac{1}{k^2\ln 2}\right\} \leq \sum_{B\subseteq[U], B\neq\emptyset, B\neq[U]} \frac{\delta}{2^{[U]}-2} = \delta.$$

□

Lemma F.4. *Let $\mathcal{F}_i$ for $i = 1\ldots$ be a filtration and $X_1, \ldots X_n$ be a sequence of Bernoulli random variables with $\mathbb{P}(X_i = 1|\mathcal{F}_{i-1}) = P_i$ with P_i being $\mathcal{F}_{i-1}$-measurable and X_i being $\mathcal{F}_i$ measurable. It holds that*

$$\mathbb{P}\left(\exists n : \sum_{t=1}^{n} X_t < \sum_{t=1}^{n} P_t/2 - W\right) \leq e^{-W}$$

Proof. $P_t - X_t$ is a Martingale difference sequence with respect to the filtration $\mathcal{F}_t$. Since X_t is nonnegative and has finite second moment, we have for any $\lambda > 0$ that $\mathbb{E}\left[e^{-\lambda(X_t - P_t)}|\mathcal{F}_{t-1}\right] \leq e^{\lambda^2 P_t/2}$ (Exercise 2.9, Boucheron et al. (2013)). Hence, we have

$$\mathbb{E}\left[e^{\lambda(P_t - X_t) - \lambda^2 P_t/2}|\mathcal{F}_{t-1}\right] \leq 1$$

and by setting $\lambda = 1$, we see that

$$M_n = e^{\sum_{t=1}^{n}(-X_t + P_t/2)}$$

is a supermartingale. It hence holds by Markov's inequality

$$\mathbb{P}\left(\sum_{t=1}^{n}(-X_t + P_t/2) \geq W\right) = \mathbb{P}\left(M_n \geq e^W\right) \leq e^{-W}\mathbb{E}[M_n] \leq e^{-W}$$

wich gives us the derised result

$$\mathbb{P}\left(\sum_{t=1}^{n} X_t \leq \sum_{t=1}^{n} P_t/2 - W\right) \leq e^{-W}$$

for a fixed n. We define now the stopping time $\tau = \min\{t \in \mathbb{N} : M_t > e^W\}$ and the sequence $\tau_n = \min\{t \in \mathbb{N} : M_t > e^W \vee t \geq n\}$. Applying the convergence theorem for nonnegative supermartingales (Theorem 5.2.9 in Durrett (2010)), we get that $\lim_{t\to\infty} M_t$ is well-defined almost surely. Therefore, M_τ is well-defined even when $\tau = \infty$. By the optional stopping theorem for nonnegative supermartingales (Theorem 5.7.6 by Durrett (2010)), we have $\mathbb{E}[M_{\tau_n}] \leq \mathbb{E}[M_0] \leq 1$ for all n and applying Fatou's lemma, we obtain $\mathbb{E}[M_\tau] = \mathbb{E}[\lim_{n\to\infty} M_{\tau_n}] \leq \liminf_{n\to\infty} \mathbb{E}[M_{\tau_n}] \leq 1$. Using Markov's inequality, we can finally bound

$$\mathbb{P}\left(\exists n : \sum_{t=1}^{n} X_t < \frac{1}{2}\sum_{t=1}^{n} P_t - W\right) \leq \mathbb{P}(\tau < \infty) \leq \mathbb{P}(M_\tau > e^W) \leq e^{-W}\mathbb{E}[M_\tau] \leq e^{-W}.$$

□